arXiv:1701.07506v1 [stat.ME] 25 Jan 2017

Bayesian Hierarchical Models with Conjugate Full-Conditional Distributions for Dependent Data from the Natural Exponential Family

Jonathan R. Bradley[1], Scott H. Holan[2], Christopher K. Wikle[2]

Abstract

We introduce a Bayesian approach for analyzing (possibly) high-dimensional dependent data that are distributed according to a member from the natural exponential family of distributions. This problem requires extensive methodological advancements, as jointly modeling high-dimensional dependent data leads to the so-called "big n problem." The computational complexity of the "big n problem" is further exacerbated when allowing for non-Gaussian data models, as is the case here. Thus, we develop new computationally efficient distribution theory for this setting. In particular, we introduce something we call the "conjugate multivariate distribution," which is motivated by the univariate distribution introduced in Diaconis and Ylvisaker (1979). Furthermore, we provide substantial theoretical and methodological development including: results regarding conditional distributions, an asymptotic relationship with the multivariate normal distribution, conjugate prior distributions, and full-conditional distributions for a Gibbs sampler. The results in this manuscript are extremely general, and can be adapted to many different settings. We demonstrate the proposed methodology through simulated examples and analyses based on estimates obtained from the US Census Bureaus' American Community Survey (ACS).

Keywords: Bayesian hierarchical model; Big data; Exponential family; Markov chain Monte Carlo; Non-Gaussian; Gibbs sampler.

[1](to whom correspondence should be addressed) Department of Statistics, Florida State University, 117 N. Woodward Ave., Tallahassee, FL 32306, bradley@stat.fsu.edu

[2]Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211-6100

1 Introduction

The multivariate normal distribution has become a fundamental tool for statisticians, as it provides a way to incorporate dependence for Gaussian and non-Gaussian data alike. Notice that many statistical models are defined hierarchically, where the joint distribution of the data, latent processes, and unknown parameters are written as the product of a data model, a latent Gaussian process model, and a parameter model (e.g., see Cressie and Wikle, 2011; Banerjee et al., 2015, among others). Jointly modeling a member from the exponential family may be seen as straightforward to some. That is, one can simply define the data model to be the appropriate member of the exponential family and define latent Gaussian processes using the hierarchical modeling framework. Models of this form are often referred to as *latent Gaussian process* (LGP) models; see Diggle et al. (1998), Rue et al. (2009), Cressie and Wikle (2011, Sections 4.1.2 and 7.1.5), and Holan and Wikle (2016), among others.

In the Bayesian context, LGPs can be nontrivial to implement using standard Markov chain Monte Carlo (MCMC) procedures when the dataset is high-dimensional. This is primarily because big data can lead to big parameter spaces, which allows parameters to be highly correlated. This in turn, creates a challenge for defining useful proposal distributions, tuning these proposal distributions, and assessing convergence of the Markov chain (e.g., see Rue et al. (2009) and Bradley et al. (2015a) for a discussion on convergence issues of MCMC algorithms for LGPs). In this article, our primary goal is to introduce new distribution theory that facilitates Bayesian inference of dependent non-Gaussian data. In particular, we introduce a multivariate distribution that leads to conjugate forms of the full conditional distributions within a Gibbs sampler.

We provide a multivariate extension of the class of distributions introduced by Diaconis and Ylvisaker (1979). The seminal paper by Diaconis and Ylvisaker (1979) developed the conjugate prior for distributions from the natural exponential family (EF), which leads to the well-known Poisson/gamma, binomial/beta, negative binomial/beta, and gamma/inverse-gamma hierarchical

models. In this article, we develop a multivariate version of this distribution, which we call the *conjugate multivariate* (CM) distribution. Similar to the special cases that emerged from Diaconis and Ylvisaker (1979), we obtain Poisson/multivariate log-gamma, binomial/multivariate logit beta, negative binomial/multivariate logit beta, and gamma/multivariate negative-inverse-gamma hierarchical models. The hierarchical model that specifies the data model to be from the natural exponential family, and the latent process to be a CM distribution is referred to as a latent CM process (LCM) model. The LCM model constitutes a more general paradigm for modeling dependent data than LGPs, since the LGP is a special case of the LCM model. An important motivating feature of this more general framework is that the LCM model incorporates dependency and results in full-conditional distributions (within a Gibbs sampler) that are easy to simulate from. This allows one to avoid computationally inefficient and subjective tuning methods.

An immediate issue that arises with the introduction of the LCM model is the need to define flexible prior distributions. One goal of this article is to describe the *fully conjugate* Bayesian hierarchical model that has a data model that belongs to the natural exponential family. By "fully conjugate" we mean that each full conditional distribution, within a Gibbs sampler, falls in the same class of distributions of the associated process or parameter models. To derive a fully conjugate statistical model, we introduce the LCM analogue to the prior distributions used in Daniels and Pourahmadi (2002), Chen and Dunson (2003), and Pourahmadi et al. (2007) for covariance parameters. Additionally, extensions of the standard inverse-gamma priors for variances of a normal random variable (Gelman, 2006) are discussed in context of the LCM.

There is an added benefit of the CM distribution besides providing conjugacy in the non-Gaussian dependent data setting. Namely, LGPs are not necessarily realistic for *every* dataset. For example, De Oliveira (2013) shows that there are parametric limitations to the LGP paradigm for count-valued data (e.g., when spatial overdispersion is small). We support this claim by showing that if certain hyperparameters (defined in Section 2) of the CM distribution are "large" then the corresponding CM distribution gives a very good approximation to a Gaussian distribution. This

indicates that if the data suggests small values of these hyperparameters, then the CM distribution should be used in place of the multivariate normal distribution.

Reduced rank methods are extremely prevalent in the more general "dependent data" setting. For example, reduced rank assumptions are crucial for principle component analysis, which has become an established technique in multivariate data analysis (e.g., see Jolliffe, 2002; Cox, 2005; Everitt and Hothorn, 2011, among others). Additionally, reduced rank models have been used to great effect within spatial and spatio-temporal settings to obtain precise predictions in a computationally efficient manner (see, e.g., Wikle and Cressie, 1999; Cressie and Johannesson, 2006; Shi and Cressie, 2007; Banerjee et al., 2008; Cressie and Johannesson, 2008; Finley et al., 2009; Katzfuss and Cressie, 2011; Cressie et al., 2010a,b; Kang and Cressie, 2011; Katzfuss and Cressie, 2012; Bradley et al., 2015b). Thus, an additional motivating feature of the LCM model is that it can easily be cast within the reduced rank modeling framework to obtain further computational gains.

This computationally efficient fully conjugate distribution theory could have an important impact on a number of different communities within and outside statistics. High-dimensional non-Gaussian data are pervasive in official statistics (e.g., see Bradley et al., 2015a), ecology (e.g., see Hooten et al., 2003; Wu et al., 2013, among others), climatology (e.g., see Wikle and Anderson, 2003), atmospheric sciences (e.g., see Sengupta et al., 2012), statistical genetics (e.g., see Lange et al., 2014, and the references therein), neuroscience (e.g., see Zhang et al., 2015; Castruccio et al., 2016), and many other domains. Despite the need for analyzing high-dimensional correlated non-Gaussian data, a majority of the Bayesian dependent data literature uses difficult to implement LGPs (Wikle and Royle, 2005; Holan and Wikle, 2012). Hence, the methodology presented here offers an exciting avenue that makes new applied research for modeling dependent non-Gaussian data practical for modern big datasets.

The LCM model is a type of the hierarchical generalized linear model (HGLM) from Lee and Nelder (1996). However, the current HGLM literature specifies an LGP for the dependent data

setting (Lee and Nelder, 2000, 2001).

There are other alternatives to a Gibbs sampler with Metropolis-Hastings updates; in particular, integrated nested Laplace approximations (INLA) (Rue et al., 2009) and Hamiltonian MCMC have proven to be useful tools in the literature. These approaches can be adapted to our new proposed distribution theory; however, the need to adapt INLA and Hamiltonian MCMC (Duane et al., 1987) to the LCM is not necessary since the full conditional distributions are straightforward to simulate from in this setting.

For Poisson counts there are a number of choices besides the LGP strategy available to incorporate dependence (e.g., see Lee and Nelder, 1974; Kotz et al., 2000; Demirhan and Hamurkaroglu, 2011, among others). For example, Wolpert and Ickstadt (1998), introduced a spatial convolution of gamma random variables, and provide a data augmentation scheme for Gibbs sampling that produces spatial predictions. The recently proposed multivariate log-gamma distribution of Bradley et al. (2015a) results in a special case of our modeling approach when the data model is Poisson, and the latent processes are are distributed according to a type of CM distribution.

The remainder of this article is organized as follows. In Section 2, we introduce the conjugate multivariate distribution and provide the necessary technical development for fully Bayesian inference of dependent data from the natural exponential family. Specifically, we define the CM distribution, give the specification of the LCM model, introduce conjugate prior distributions, and derive the full conditional distributions for the Gibbs sampler. We end Section 2 with a brief summary of the LCM model used in subsequent sections. Then, in Section 3 we provide a simulation study to show the performance of the LCM model compared to LGPs. In Section 4, a demonstration is presented using estimates obtained from the American Community Survey (ACS). Finally, Section 5 contains discussion. For convenience of exposition, proofs of the technical results, and instructions on implementation are given in the Supplemental Appendix.

2 Distribution Theory for Dependent Data from the Natural Exponential Family

In this section, we propose methodology for Bayesian analysis of non-Gaussian dependent data from the natural exponential family. In Section 2.1, we review and develop the univariate distribution introduced in Diaconis and Ylvisaker (1979). Then, in Section 2.2, this univariate distribution is used as the rudimentary quantity to develop the CM distribution. This new multivariate distribution theory is incorporated within a Bayesian hierarchical model (i.e., the aforementioned LCM model) in Section 2.3. Prior distributions for unknown parameters of the LCM model are discussed in Sections 2.4 and 2.5. The information in Sections 2.1 through 2.5 are consolidated into a brief model summary in Section 2.6.

2.1 The Diaconis and Ylvisaker Conjugate Distribution

Suppose Z is distributed according to the natural exponential family (Diaconis and Ylvisaker, 1979; Lehmann and Casella, 1998), then

$$f(\mathbf{Z}|Y) = \exp\{ZY - b\psi(Y) + c(Z)\};\ Z \in \mathscr{Z}, Y \in \mathscr{Y}, \tag{1}$$

where f will be used to denote a generic probability density function/probability mass function (pdf/pmf), $Z \in \mathscr{Z}$, $\mathscr{Z}$ is the support of Z, $\mathscr{Y}$ is the support of Y, b is known, and both $\psi(\cdot)$ and $c(\cdot)$ are known real-valued functions. The function $b\psi(Y)$ is often called the log partition function (Lehmann and Casella, 1998). It will be useful for us to discuss $\psi(Y)$ and not $b\psi(Y)$; hence, we refer to $\psi(Y)$ as the "unit log partition function" because it's coefficient is one and not b. Let $\text{EF}(Y;\ \psi)$ denote a shorthand for the pdf in (1). It follows from Diaconis and Ylvisaker (1979) that

the conjugate prior distribution for Y is given by,

$$f(Y|\alpha,\kappa) = K(\alpha,\kappa)\exp\{\alpha Y - \kappa\psi(Y)\};\ Y \in \mathscr{Y}, \frac{\alpha}{\kappa} \in \mathscr{Z}, \kappa > 0, \tag{2}$$

where $K(\alpha,\kappa)$ is a normalizing constant. Let $\mathrm{DY}(\alpha,\kappa;\ \psi)$ denote a shorthand for the pdf in (2). Here "DY" stands for "Diaconis-Ylvisaker," and we will refer to Y as either a Diaconis-Ylvisaker random variable or a DY random variable. Diaconis and Ylvisaker (1979) proved that the pdf in (2) is proper (i.e., yields a probability measure). We also call α and κ "DY parameters." For examples of ψ, $\mathrm{EF}(Y;\psi)$, and $K(\alpha,\kappa)$ see Table 1.

Multiplying both sides of (2) by $\exp(tY)$ and integrating, gives the moment generating function

$$E[\exp(tY)|\alpha,\kappa] = \frac{K(\alpha,\kappa)}{K(\alpha+t,\kappa)}, \tag{3}$$

which exists provided that $(\alpha+t)/\kappa \in \mathscr{Z}$, $\kappa > 0$, and the corresponding values of $K(\alpha+t,\kappa)$ and $K(\alpha,\kappa)$ are strictly positive and finite. This gives us that the mean and variance of Y is

$$E(Y|\alpha,\kappa) = K(\alpha,\kappa)K^{(1)}(\alpha,\kappa) \tag{4}$$

$$\mathrm{var}(Y|\alpha,\kappa) = K(\alpha,\kappa)K^{(2)}(\alpha,\kappa) - K(\alpha,\kappa)^2K^{(1)}(\alpha,\kappa)^2, \tag{5}$$

assuming that the moment generating function exists at $t=0$, where $K^{(1)}(\alpha,\kappa) \equiv \left[\frac{d}{dt}\frac{1}{K(\alpha+t,\kappa)}\right]_{t=0}$ and $K^{(2)}(\alpha,\kappa) \equiv \left[\frac{d^2}{dt^2}\frac{1}{K(\alpha+t,\kappa)}\right]_{t=0}$.

Finally, it is immediate from (1) and (2) that

$$Y|Z,\alpha,\kappa \sim \mathrm{DY}\left(\alpha+Z,\kappa+b;\ \psi\right). \tag{6}$$

This conjugacy motivates the development of a multivariate version of the DY random variable to model dependent non-Gaussian data from the natural exponential family. Thus, in this section, we

Data Model	**Natural Parameter**	**Log Partition Function** (i.e., ψ and b)	**Normalizing Constant**	**How to Simulate From the DY Distribution**
Gamma(a,k) $f(Z\|\alpha,\kappa)=\frac{1}{\Gamma(a)k^a}\exp(-Z/k)$ $a>0,k>0,Z>0$	Negative Reciprocal: $Y=-\frac{1}{k}$.	$\psi_1(Y)=\log\left(-\frac{1}{Y}\right)$ $b=a$	$K(\alpha,\kappa)=\frac{1}{\Gamma(\kappa+1)\alpha^{\kappa+1}}$	Let $W\sim$ Gamma$(\kappa+1,1/\alpha)$, where $\alpha>0$, and $\kappa>0$. Then, $-W\sim$ DY$(\alpha,\kappa;\psi_1)$.
Bin(t,p) $f(Z\|t,p)=\binom{t}{Z}p^Z(1-p)^{t-Z}$ $0<p<1,t=1,2,\ldots,Z=0,\ldots,t$	Logit: $Y=\log\left(\frac{p}{1-p}\right)$	$\psi_2(Y)=\log(1+\exp(Y))$ $b=t$	$K(\alpha,\kappa)=\frac{\Gamma(\kappa)}{\Gamma(\alpha)\Gamma(\kappa-\alpha)}$	Let $W\sim$ Beta$(\alpha,\kappa-\alpha)$, where $\kappa>\alpha>0$ and "Beta$(\alpha,\kappa-\alpha)$" is a shorthand for the beta distribution with shape parameter α and scale parameter $\kappa-\alpha$. Then, $\log\left(\frac{W}{1-W}\right)\sim$ DY$(\alpha,\kappa;\psi_2)$.
NegBin(t,p) $f(Z\|t,p)=\binom{Z+t-1}{Z}p^Z(1-p)^t$ $0\le p\le 1,t=1,2,\ldots,Z=0,1,\ldots,$	Logit: $Y=\log\left(\frac{p}{1-p}\right)$	$\psi_2(Y)=\log(1+\exp(Y))$ $b=t+Z$	$K(\alpha,\kappa)=\frac{\Gamma(\kappa)}{\Gamma(\alpha)\Gamma(\kappa-\alpha)}$	Let $W\sim$ Beta$(\alpha,\kappa-\alpha)$, where $\kappa>\alpha>0$. Then, $\log\left(\frac{W}{1-W}\right)\sim$ DY$(\alpha,\kappa;\psi_2)$.
Pois(μ) $f(Z\|\mu)=\frac{\mu^Z\exp(-\mu)}{Z!}$ $\mu\in\mathbb{R}^+,Z=0,1,2,\ldots$	Log $Y=\log(\mu)$	$\psi_3(Y)=\exp(Y)$ $b=1$	$K(\alpha,\kappa)=\frac{1}{\Gamma(\alpha)\kappa^\alpha}$	Let $W\sim$ Gamma$(\alpha,1/\kappa)$, where $\alpha>0$ and $\kappa>0$. Then, $\log(W)\sim$ DY$(\alpha,\kappa;\psi_3)$.
Norm(μ,s) $f(Z\|\mu,s)=\left(\frac{1}{2\pi s^2}\right)^{1/2}\exp\left(-\frac{-(Z-\mu)^2}{s^2}\right)$ $\mu\in\mathbb{R},s\in\mathbb{R}^+,Z\in\mathbb{R}$	Linear: $Y=\frac{\mu}{s}$	$\psi_4(Y)=Y^2$ $b=\frac{1}{2}$	$K(\alpha,\kappa)=\left(\frac{\kappa}{\pi}\right)^{1/2}\exp(-\frac{\alpha^2}{4\kappa})$	Let W be a normal random variable with mean $\frac{\alpha}{2\kappa}$ and variance $\frac{1}{2\kappa}$. Then, $W\sim$ DY$(\alpha,\kappa;\psi_4)$.

Table 1: Univariate Distributions: The first column has the data model, the second column has the natural parameter, the third column contains quantities that define the log partition function, the fourth column has the normalizing constant, and the fifth column has instructions on how to simulate from the DY random variable with the corresponding ψ. Let $\mathbb{R}^+=\{x:x<0\}$.

define a conjugate multivariate distribution and develop a distribution theory that will prove to be extremely useful for fully Bayesian analysis in the dependent non-Gaussian (natural exponential family) data setting.

2.2 The Conjugate Multivariate Distribution

Unit Log Partition Function (i.e., ψ)	**CM Distribution (i.e.,** $f(\mathbf{Y}\|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})$**)**
$\psi_1(Y) = \log\left(-\frac{1}{Y}\right)$	$\det(\mathbf{V}^{-1})\left\{\prod_{i=1}^n \frac{1}{\Gamma(\kappa_i+1)\alpha_i^{\kappa_i+1}}\right\}\exp\left[\boldsymbol{\alpha}'\mathbf{s}-\boldsymbol{\kappa}'\log\left(-\mathbf{s}^{(-1)}\right)\right] I(-\mathbf{s}\in\mathbb{R}_n^+)$
$\psi_2(Y) = \log(1+\exp(Y))$	$\det(\mathbf{V}^{-1})\left\{\prod_{i=1}^n \frac{\Gamma(\kappa_i)}{\Gamma(\alpha_i)\Gamma(\kappa_i-\alpha_i)}\right\}\exp\left[\boldsymbol{\alpha}'\mathbf{V}^{-1}(\mathbf{Y}-\boldsymbol{\mu})-\boldsymbol{\kappa}'\log\left[\mathbf{J}_{n,1}+\exp\left\{\mathbf{V}^{-1}(\mathbf{Y}-\boldsymbol{\mu})\right\}\right]\right] I(\mathbf{Y}\in\mathbb{R}^n)$
$\psi_3(Y) = \exp(Y)$	$\det(\mathbf{V}^{-1})\left\{\prod_{i=1}^n \frac{1}{\Gamma(\alpha_i)\kappa_i^{\alpha_i}}\right\}\exp\left[\boldsymbol{\alpha}'\mathbf{V}^{-1}(\mathbf{Y}-\boldsymbol{\mu})-\boldsymbol{\kappa}'\exp\left\{\mathbf{V}^{-1}(\mathbf{Y}-\boldsymbol{\mu})\right\}\right] I(\mathbf{Y}\in\mathbb{R}^n)$
$\psi_4(Y) = Y^2$	$\det(\mathbf{V}^{-1})\left\{\prod_{i=1}^n \left(\frac{\kappa_i}{\pi}\right)^{1/2}\right\}\exp\left\{-(\mathbf{Y}-\boldsymbol{\mu}-\boldsymbol{\gamma})'\mathbf{V}^{-1}\boldsymbol{\Sigma}^{-1}\mathbf{V}^{-1\prime}(\mathbf{Y}-\boldsymbol{\mu}-\boldsymbol{\gamma})/2\right\} I(\mathbf{Y}\in\mathbb{R}^n)$

Table 2: Special Cases: We list the form of the CM distribution by ψ_j for $j = 1,\ldots,4$. The first column has the unit log partition function ψ_j, and the second column has the form of the CM distribution with generic $\mathbf{V}^{-1}\in\mathbb{R}^n\times\mathbb{R}^n$. Let $\mathbf{J}_{m,g}$ denote a $m\times g$ matrix of ones, $\mathbf{s} = (s_1,\ldots,s_n)' \equiv \mathbf{V}^{-1}(\mathbf{Y}-\boldsymbol{\mu})$, $\mathbf{s}^{(-1)} = (1/s_1,\ldots,1/s_n)'$, $\boldsymbol{\gamma} = \left(\frac{\alpha_1}{2\kappa_1},\ldots,\frac{\alpha_n}{2\kappa_n}\right)'$, and $\boldsymbol{\Sigma}\equiv \text{diag}\left(\frac{1}{2\kappa_i} : i = 1,\ldots,n\right)$.

Bradley et al. (2015a) use a linear combination of independent log-gamma random variables to build their multivariate log-gamma distribution. In a similar manner we take linear combinations of DY random variables to generate a conjugate version of the DY distribution. Specifically, let the n-dimensional random vector $\mathbf{w} = (w_1,\ldots,w_n)'$ consist of n mutually independent DY random variables such that $w_i \sim \text{DY}(\alpha_i,\kappa_i;\,\psi)$ for $i = 1,\ldots,n$. Then, define $\mathbf{Y}\equiv(Y_1,\ldots,Y_n)'$ such that

$$\mathbf{Y} = \boldsymbol{\mu} + \mathbf{V}\mathbf{w}, \tag{7}$$

where $\mathbf{Y}\in\mathscr{M}^n$, the matrix $\mathbf{V}\in\mathbb{R}^n\times\mathbb{R}^n$, and $\boldsymbol{\mu}\in\mathbb{R}^n$. The space $\mathscr{M}^n$ is not necessarily equal to $\mathscr{Y}^n \equiv \{\mathbf{Y} = (Y_1,\ldots,Y_n)' : Y_i\in\mathscr{Y}, i = 1,\ldots,n\}$; for example, if $\mathscr{M}$ is strictly positive, we obtain a $\mathbf{Y}$ that is negative since $\mathbf{V}\in\mathbb{R}^n\times\mathbb{R}^n$. Call $\mathbf{Y}$ in (7) a *conjugate multivariate* (CM) random

vector. A special case of the CM random vector is the multivariate normal random vector. To see this, let $\alpha_i \equiv 0$, $\kappa_i \equiv 1/2$, and $\psi(Y) = Y^2$ for $Y \in \mathbb{R}$. Then, it follows that (7) is a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{V}\mathbf{V}'$, since the elements of $\mathbf{w}$ consist of i.i.d. standard normal random variables. Additionally, the aforementioned MLG distribution can be written as a CM distribution when $\alpha > 0$, $\kappa > 0$, and $\psi(Y) = \exp(Y)$.

To use the CM distribution in a Bayesian context, we require its pdf, which is formally stated below.

Proposition 1: Let $\boldsymbol{Y} = \boldsymbol{\mu} + \boldsymbol{V}\boldsymbol{w}$*, where* $\boldsymbol{Y} = (Y_1, \ldots, Y_n)'$*,* $\boldsymbol{\mu} \in \mathbb{R}^n$*, the* $n \times n$ *real valued matrix* $\boldsymbol{V}$ *is invertible, and the n-dimensional random vector* $\boldsymbol{w} = (w_1, \ldots, w_n)'$ *consists of n mutually independent DY random variables such that* $w_i \sim \mathrm{DY}(\alpha_i, \kappa_i;\ \psi)$ *for* $i = 1, \ldots, n$*.*

(i) Then $\boldsymbol{Y}$ *has the following pdf:*

$$f(\boldsymbol{Y}|\boldsymbol{\mu}, \boldsymbol{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = \det(\boldsymbol{V}^{-1}) \left\{ \prod_{i=1}^{n} K(\kappa_i, \alpha_i) \right\} \exp\left[\boldsymbol{\alpha}'\boldsymbol{V}^{-1}(\boldsymbol{Y} - \boldsymbol{\mu}) - \boldsymbol{\kappa}'\psi\left\{ \boldsymbol{V}^{-1}(\boldsymbol{Y} - \boldsymbol{\mu}) \right\} \right] I(\boldsymbol{Y} \in \mathscr{M}^n), \tag{8}$$

where $I(\cdot)$ *is the indicator function, the j-th element of* $\psi\left\{\boldsymbol{V}^{-1}(\boldsymbol{Y} - \boldsymbol{\mu})\right\}$ *contains* ψ *evaluated at the j-th element of the n-dimensional vector* $\boldsymbol{V}^{-1}(\boldsymbol{Y} - \boldsymbol{\mu})$*, "det" denotes the determinant function,* $\boldsymbol{\alpha} \equiv (\alpha_1, \ldots, \alpha_n)'$*, and* $\boldsymbol{\kappa} \equiv (\kappa_1, \ldots, \kappa_n)'$*.*

(ii) The mean and variance of $\boldsymbol{Y}$ *is given by,*

$$\begin{aligned} E(\boldsymbol{Y}|\boldsymbol{\alpha}, \boldsymbol{\kappa}) &= \boldsymbol{\mu} + \boldsymbol{V}\boldsymbol{k}(\boldsymbol{\alpha}, \boldsymbol{\kappa}) \\ \mathrm{cov}(\boldsymbol{Y}|\boldsymbol{\alpha}, \boldsymbol{\kappa}) &= \boldsymbol{V}\boldsymbol{K}(\boldsymbol{\alpha}, \boldsymbol{\kappa})\boldsymbol{V}', \end{aligned} \tag{9}$$

where, the n-dimensional real-valued vector

$$\boldsymbol{k}(\boldsymbol{\alpha},\boldsymbol{\kappa}) = \Big(K(\alpha_1,\kappa_1)K^{(1)}(\alpha_1,\kappa_1),\ldots,K(\alpha_n,\kappa_n)K^{(1)}(\alpha_n,\kappa_n)\Big)',$$

and the $n\times n$ *diagonal matrix* $\boldsymbol{K}(\boldsymbol{\alpha},\boldsymbol{\kappa}) \equiv \text{diag}\left\{K(\alpha_i,\kappa_i)K^{(2)}(\alpha_i,\kappa_i) - K(\alpha_i,\kappa_i)^2K^{(1)}(\alpha_i,\kappa_i)^2\right\}$.

The proof of Proposition 1(i) can be found in the Supplemental Appendix, and in general, let CM($\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa};\ \psi$) be shorthand for the pdf in (8). Proposition 1(ii) follows immediately from Equations (4) and (5), and thus, Proposition 1(ii) is stated without proof.

When comparing (1), (2), and (8) we see that the univariate natural exponential family, the DY pdf, and the CM pdf share a basic structure. Specifically, all three distributions have an exponential term and an "exponential of $-\psi$ term." This pattern is the main reason why conjugacy exists between the distributions from the natural exponential family and the DY distribution, which we take advantage of in subsequent sections. Also, Proposition $1(ii)$, shows that if we restrict $\mathbf{V}$ (or equivalently $\mathbf{V}^{-1}$) to be a lower unit triangle matrix, then the expression of the covariance matrix of $\mathbf{Y}$ in (9) is a type of LDL decomposition (Ravishanker and Dey, 2002). Hence, in subsequent sections we assume that $\mathbf{V}$ is lower unit triangular.

In this article, we consider CM distributions that are implied by the unit log partition function of the data model including: the gamma data model, binomial data model, negative binomial data model, the Poisson data model, and the normal data model (see Tables 1 and 3). In the univariate case, each of these special cases lead to well-known hierarchical models (i.e., gamma/inverse-gamma, (negative) binomial/beta, Poisson/log-gamma, and normal/normal models) (Diaconis and Ylvisaker, 1979). To delineate from the univariate setting, we shall refer to CM($\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa};\psi_j$) for $j = 1,\ldots,4$ (see Table 1 for the definitions of $\psi_1,\psi_2,\ \psi_3$, and ψ_4) as the multivariate negative-inverse-gamma distribution, multivariate logit-beta distribution, the multivariate log-gamma, and the multivariate normal distribution, respectively.

These choices of the CM distribution are themselves extremely general. For example, when $\boldsymbol{\alpha} = \mathbf{J}_{n,1}$, we obtain an exponential/multivariate negative-inverse-gamma model. Similarly, the binomial/multivariate logit-beta model has a Bernoulli/multivariate logit-beta model as a special case, which occurs when the number of Bernoulli trials that define the binomial distribution is equal to one. Likewise, when the number of successful Bernoulli trials is equal to one, the negative binomial/multivariate logit-beta model reduces to a geometric/multivariate logit-beta specification. This creates opportunity for analyzing many different types of dependent data.

Bayesian inference not only requires the pdf of $\mathbf{Y}$, but also requires simulating from conditional distributions of $\mathbf{Y}$. Thus, we provide the technical results needed to simulate from these conditional distributions.

Theorem 1: Let $\boldsymbol{Y} \sim \mathrm{CM}(\boldsymbol{\mu}, \boldsymbol{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa};\ \psi)$*, and let* $\boldsymbol{Y} = (Y_1, \ldots, Y_n)' = (\boldsymbol{Y}_1', \boldsymbol{Y}_2')'$*, so that* $\boldsymbol{Y}_1$ *is* r*-dimensional and* $\boldsymbol{Y}_2$ *is* $(n-r)$*-dimensional. In a similar manner, partition* $\boldsymbol{V}^{-1} = [\boldsymbol{H}\ \boldsymbol{B}]$ *into an* $n \times r$ *matrix* $\boldsymbol{H}$ *and an* $n \times (n-r)$ *matrix* $\boldsymbol{B}$*. Also let* $\boldsymbol{\mu}^* = \boldsymbol{V}^{-1}\boldsymbol{\mu} - \boldsymbol{B}\boldsymbol{d}$ *for* $\boldsymbol{d} \in \mathbb{R}^{n-r}$*. Then, the conditional distribution* $\boldsymbol{Y}_1 | \boldsymbol{Y}_2 = \boldsymbol{d}, \boldsymbol{\mu}, \boldsymbol{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}$ *is given by*

$$f(\boldsymbol{Y}_1 | \boldsymbol{Y}_2 = \boldsymbol{d}, \boldsymbol{\mu}, \boldsymbol{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = M\ \exp\left\{\boldsymbol{\alpha}'\boldsymbol{H}\boldsymbol{Y}_1 - \boldsymbol{\kappa}'\psi(\boldsymbol{H}\boldsymbol{Y}_1 - \boldsymbol{\mu}^*)\right\} I\{(\boldsymbol{Y}_1', \boldsymbol{d}')' \in \mathscr{M}^n\}, \tag{10}$$

where M *is a strictly positive and finite normalizing constant. Let* $\mathrm{CM_c}(\boldsymbol{\mu}^*, \boldsymbol{H}, \boldsymbol{\alpha}, \boldsymbol{\kappa};\ \psi)$ *be a short-hand for the pdf in (10), where the subscript "c" represents the word "conditional."*

It is not immediately clear how to simulate from a $\mathrm{CM_c}$ distribution since $\mathbf{H}$ in (10) is not square, and hence, one can not use Equation (7) directly. Thus, we require an additional result that allows us to simulate from $\mathrm{CM_c}$.

Theorem 2: Let $\boldsymbol{w} \sim \mathrm{CM}\,(\boldsymbol{\mu}, \boldsymbol{I}_n, \boldsymbol{\alpha}, \boldsymbol{\kappa};\ \psi)$*, where* $\boldsymbol{\mu} \in \mathbb{R}^n$ *and* $\boldsymbol{I}_n$ *is an* $n \times n$ *identity matrix. Assume*

that the function ψ is absolutely continuous, and define $\boldsymbol{H} \in \mathbb{R}^n \times \mathbb{R}^r$.

(i) Then the r-dimensional random vector $\boldsymbol{q} = (q_1, \ldots, q_r)' \equiv (\boldsymbol{H}'\boldsymbol{H})^{-1}\boldsymbol{H}'\boldsymbol{w}$ has the probability density function

$$f(\boldsymbol{q}|\boldsymbol{\mu}, \boldsymbol{H}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = M_1 \ \exp\left\{\boldsymbol{\alpha}'\boldsymbol{H}\boldsymbol{q} - \boldsymbol{\kappa}'\psi(\boldsymbol{H}\boldsymbol{q} - \boldsymbol{\mu})\right\}\{I(\boldsymbol{q} \in \mathscr{M}_r^n)\} = \mathrm{CM_c}(\boldsymbol{\mu}, \boldsymbol{H}, \boldsymbol{\alpha}, \boldsymbol{\kappa};\ \psi), \tag{11}$$

where $\mathscr{M}_r^n \equiv \{\boldsymbol{q} : \boldsymbol{Y} \in \mathscr{M}^n, \boldsymbol{q} = (\boldsymbol{I}_r, \boldsymbol{0}_{r,n-r})\boldsymbol{Y}\}$, $\boldsymbol{0}_{m,g}$ is a $m \times g$ matrix of zeros, and M_1 is a strictly positive and finite normalizing constant.

(ii) The conditional mean and covariance can be computed as

$$E(\boldsymbol{Y}_1|\boldsymbol{Y}_2 = \boldsymbol{d}, \boldsymbol{\mu}, \boldsymbol{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = (\boldsymbol{H}'\boldsymbol{H})^{-1}\boldsymbol{H}'\{\boldsymbol{\mu}^* + \boldsymbol{k}(\boldsymbol{\alpha}, \boldsymbol{\kappa})\} \tag{12}$$

$$\mathrm{cov}(\boldsymbol{Y}_1|\boldsymbol{Y}_2 = \boldsymbol{d}, \boldsymbol{\mu}, \boldsymbol{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = (\boldsymbol{H}'\boldsymbol{H})^{-1}\boldsymbol{H}'\boldsymbol{K}(\boldsymbol{\alpha}, \boldsymbol{\kappa})\boldsymbol{H}(\boldsymbol{H}'\boldsymbol{H})^{-1}, \tag{13}$$

where recall that $\boldsymbol{\mu}^*$ is defined in the statement of Theorem 1.

The proof of Theorem $2(i)$ is given in the Supplemental Appendix. The proof of Theorem $2(ii)$ follows immediately from Proposition $1(ii)$ and Theorem $2(i)$. Thus, we state Theorem $2(ii)$ without proof.

Theorem 2 shows that it is (computationally) easy to simulate from $\mathrm{CM_c}$ provided that $r \ll n$. Recall that $\mathbf{H}$ is $n \times r$, which implies that computing the $r \times r$ matrix $(\mathbf{H}'\mathbf{H})^{-1}$ is computationally feasible when r is "small." Furthermore, Theorem 2 shows that the r-dimensional random vector $\mathbf{q}$ is an orthogonal projection of the n-dimensional random vector $\mathbf{w}$ onto the column space spanned by the columns of $\mathbf{H}$. This provides a geometric interpretation of random vectors generated from $\mathrm{CM_c}$.

2.3 The LCM Model

The LCM model can be written as the product of the following conditional and marginal distributions:

$$
\begin{aligned}
&\text{Data Model}: \; Z_i|\boldsymbol{\beta},\boldsymbol{\eta},\xi_i \overset{\text{ind}}{\sim} \text{EF}\left(\mathbf{x}_i'\boldsymbol{\beta}+\boldsymbol{\phi}_i'\boldsymbol{\eta}+\xi_i;\, \psi_j\right); \quad i=1,\ldots,n, j=1,\ldots,4\\
&\text{Process Model 1}: \; \boldsymbol{\eta}|\mathbf{V},\alpha_\eta,\kappa_\eta \sim \text{CM}\left(\mathbf{0}_{r,1},\mathbf{V},\boldsymbol{\alpha}_\eta,\boldsymbol{\kappa}_\eta;\, \psi_k\right);\\
&\text{Process Model 2}: \; \boldsymbol{\xi}|\boldsymbol{\alpha}_\xi,\boldsymbol{\kappa}_\xi \sim \text{CM}\left(\mathbf{0}_{n,1},\mathbf{V}_\xi,\boldsymbol{\alpha}_\xi,\boldsymbol{\kappa}_\xi;\, \psi_k\right);\\
&\text{Parameter Model 1}: \; \boldsymbol{\beta}|\alpha_\beta,\kappa_\beta \sim \text{CM}\left(\mathbf{0}_{p,1},\mathbf{V}_\beta,\boldsymbol{\alpha}_\beta,\boldsymbol{\kappa}_\beta;\, \psi_k\right); k=1,\ldots,4,
\end{aligned}
\tag{14}
$$

where ψ_j and ψ_k (for $j,k=1,\ldots,4$) are defined in Table 1 and the elements of n-dimensional vector $\mathbf{Z} \equiv (Z_1,\ldots,Z_n)'$ represent data that can be reasonably modeled using a member from the natural exponential family. Additionally for each i, $\mathbf{x}_i$ is a known p-dimensional vector of covariates, $\boldsymbol{\beta}=(\beta_1,\ldots,\beta_p)' \in \mathbb{R}^p$ is an unknown vector interpreted as fixed effects, $\boldsymbol{\phi}_i$ is a known r-dimensional real-valued vector (see Section 4.2 for an example), and the r-dimensional vector $\boldsymbol{\eta}=(\eta_1,\ldots,\eta_r)'$ and n-dimensional vector $\boldsymbol{\xi} \equiv (\xi_1,\ldots,\xi_n)'$ are interpreted as real-valued random effects. We have not yet provided specifications of the hyperparameters and variance parameters: $\boldsymbol{\alpha}_\beta=(\alpha_{\beta,1},\ldots,\alpha_{\beta,p})'$, $\boldsymbol{\alpha}_\eta=(\alpha_{\eta,1},\ldots,\alpha_{\eta,r})'$, $\boldsymbol{\alpha}_\xi=(\alpha_{\xi,1},\ldots,\alpha_{\xi,n})'$, $\boldsymbol{\kappa}_\beta=(\kappa_{\beta,1},\ldots,\kappa_{\beta,p})'$, $\boldsymbol{\kappa}_\eta=(\kappa_{\eta,1},\ldots,\kappa_{\eta,r})'$, $\boldsymbol{\kappa}_\xi=(\kappa_{\xi,1},\ldots,\kappa_{\xi,n})'$, $\mathbf{V}_\beta \in \mathbb{R}^p \times \mathbb{R}^p$, $\mathbf{V} \in \mathbb{R}^r \times \mathbb{R}^r$, and $\mathbf{V}_\xi \in \mathbb{R}^n \times \mathbb{R}^n$, where $\alpha_{\beta,i}/\kappa_{\beta,i} \in \mathscr{Y}$, $\alpha_{\eta,j}/\kappa_{\eta,j} \in \mathscr{Y}$, $\alpha_{\xi,k}/\kappa_{\xi,k} \in \mathscr{Y}$, $\kappa_{\beta,i}>0$, $\kappa_{\eta,j}>0$, and $\kappa_{\xi,k}>0$; $i=1,\ldots,p$, $j=1,\ldots,r$, $k=1,\ldots,n$. These details are presented in Sections 2.4 and 2.5.

An important point argued in the Section 1 is that the LGP model is a *special case* of an LCM. This can now easily be seen by letting $j=1,\ldots,4$, $k=4$, $\boldsymbol{\alpha}_\beta=\mathbf{0}_{p,1}$, $\boldsymbol{\alpha}_\eta=\mathbf{0}_{r,1}$, and $\boldsymbol{\alpha}_n=\mathbf{0}_{n,1}$. This specification yields an LGP model. In this article, we consider setting $j=k$, since this leads to expressions of the full conditional distributions for $\boldsymbol{\beta}$, $\boldsymbol{\eta}$, and $\boldsymbol{\xi}$ that are easy to simulate from.

Theorem 3: Suppose the n-dimensional data vector $\boldsymbol{Z}$ *is distributed according to the model in (14) with* $j = k$ *and* $j = 1,\ldots,4$*. We have the following full conditional distributions for* $\boldsymbol{\beta}$*,* $\boldsymbol{\eta}$*, and* $\boldsymbol{\xi}$*.*

$$
\begin{aligned}
f(\boldsymbol{\beta}|\cdot) &= \mathrm{CM_c}(\boldsymbol{\mu}_\beta, \boldsymbol{H}_\beta, \boldsymbol{\alpha}^*_\beta, \boldsymbol{\kappa}^*_\beta;\, \psi_j)\left\{\prod_{i=1}^{n} I(\boldsymbol{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\right\} \\
f(\boldsymbol{\eta}|\cdot) &= \mathrm{CM_c}(\boldsymbol{\mu}_\eta, \boldsymbol{H}_\eta, \boldsymbol{\alpha}^*_\eta, \boldsymbol{\kappa}^*_\eta;\, \psi_j)\left\{\prod_{i=1}^{n} I(\boldsymbol{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\right\} \\
f(\boldsymbol{\xi}|\cdot) &= \mathrm{CM_c}(\boldsymbol{\mu}_\xi, \boldsymbol{H}_\xi, \boldsymbol{\alpha}^*_\xi, \boldsymbol{\kappa}^*_\xi;\, \psi_j)\left\{\prod_{i=1}^{n} I(\boldsymbol{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\right\},
\end{aligned} \tag{15}
$$

where $f(\boldsymbol{\beta}|\cdot)$ *represents the pdf of* $\boldsymbol{\beta}$ *given all other process variables, parameters, and the data. We define* $f(\boldsymbol{\eta}|\cdot)$ *and* $f(\boldsymbol{\xi}|\cdot)$ *in a similar manner. For ease of exposition, the definitions of the remaining unknown quantities in (15) are provided in Table 3.*

For ψ_2, ψ_3, and ψ_4 we have that $\mathscr{Y} = \mathbb{R}$, $\mathscr{M}_p^{n+p} = \mathbb{R}^p$, $\mathscr{M}_r^{n+r} = \mathbb{R}^r$, and $\mathscr{M}_n^{2n} = \mathbb{R}^n$. This implies that if $\boldsymbol{\beta} \in \mathscr{M}_p^{n+p} = \mathbb{R}^p$, $\boldsymbol{\eta} \in \mathscr{M}_r^{n+r} = \mathbb{R}^r$, and $\boldsymbol{\xi} \in \mathscr{M}_n^{2n} = \mathbb{R}^n$ then $\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y} = \mathbb{R}$, since for $i = 1,\ldots,n$, $\mathbf{x}_i$ and $\boldsymbol{\psi}_i$ are real-valued. Consequently, for $j = 2,3$, and 4, we can replace (15) with

$$
\begin{aligned}
f(\boldsymbol{\beta}|\cdot) &= \mathrm{CM_c}(\boldsymbol{\mu}_\beta, \mathbf{H}_\beta, \boldsymbol{\alpha}^*_\beta, \boldsymbol{\kappa}^*_\beta;\, \psi_j) \\
f(\boldsymbol{\eta}|\cdot) &= \mathrm{CM_c}(\boldsymbol{\mu}_\eta, \mathbf{H}_\eta, \boldsymbol{\alpha}^*_\eta, \boldsymbol{\kappa}^*_\eta;\, \psi_j) \\
f(\boldsymbol{\xi}|\cdot) &= \mathrm{CM_c}(\boldsymbol{\mu}_\xi, \mathbf{H}_\xi, \boldsymbol{\alpha}^*_\xi, \boldsymbol{\kappa}^*_\xi;\, \psi_j),
\end{aligned}
$$

which is advantageous because one can use Theorem $2(i)$ directly to sample from the full conditional-distributions. Furthermore, if $p \ll n$ and $r \ll n$ it is computationally easy to simulate from $f(\boldsymbol{\beta}|\cdot)$ and $f(\boldsymbol{\eta}|\cdot)$, since Theorem $2(i)$ requires one to compute $(\mathbf{H}_\beta'\mathbf{H}_\beta)^{-1}$ and $(\mathbf{H}_\eta'\mathbf{H}_\eta)^{-1}$ and these computations are on the order of p^3 and r^3, respectively. Additionally, since $\mathbf{H}_\xi$ implies a diagonal matrix for $\mathbf{H}_\xi'\mathbf{H}_\xi$ it is also straightforward to simulated from $f(\boldsymbol{\xi}|\cdot)$ using Theorem $2(i)$.

The impact of having easy to sample full conditional distributions is more easily seen when one considers the alternative (i.e., the LGP specification of the LCM). Simulating from the analogous full-conditional distributions of an LGP is something that typically requires ad-hoc Markov chain Monte Carlo (MCMC) techniques (e.g., see the discussions in Rue et al., 2009; Bradley et al., 2015a), and this presents an added barrier for Bayesian methodology from the point-of-view of reproducible research. Thus, the implications of Theorem 3 are important for Bayesian analysis because it is now straightforward to update fixed and random effects within a Gibbs sampler for many more classes of non-Gaussian dependent data.

2.4 Prior Distributions on Covariance Parameters

A critical feature of our proposed distribution theory is the incorporation of dependence in non-Gaussian data from the exponential family. From this point-of-view it is especially important to learn about these dependencies, which are quantified by the unknown $n \times n$ real-valued matrix $\mathbf{V}$ (or equivalently $\mathbf{V}^{-1}$). Thus, we place a prior distribution on $\mathbf{V}^{-1}$. Specifically, let $\mathbf{V}^{-1}$ be an unknown lower unit triangle matrix. That is, let $\mathbf{V}^{-1} \equiv \left\{v_{i,j}\right\}$, where $v_{i,j} = 1$ for $j = i$, $v_{i,j} = 0$ for $j > i$, and $v_{i,j} \in \mathbb{R}$ for $j < i$. It will be useful to organize the elements below the lower main diagonal into the $(i-1)$-dimensional vectors $\mathbf{v}_i \equiv \left(v_{i,j} : j = 1, \ldots, i-1\right)'$ for $i = 2, \ldots, n$.

We place a CM prior distribution on $\mathbf{v}_i$ for each i. Specifically, let

$$\mathbf{v}_i \overset{\text{ind}}{\sim} \text{CM}(\mathbf{0}, \mathbf{C}_i, \boldsymbol{\alpha}_i, \boldsymbol{\kappa}_i;\ \psi);\ i = 2, \ldots, n, \tag{16}$$

where, in practice, the $(i-1) \times (i-1)$ matrix $\mathbf{C}_i$ is set equal to $\sigma_v \mathbf{I}_{i-1}$, and $\boldsymbol{\alpha}_i$, $\boldsymbol{\kappa}_i$, and σ_v are specified such that (16) is relatively "flat" (see Section 2.6 for more discussion). This specification leads to a conjugate full-conditional distribution within a Gibbs sampler.

Theorem 4: Suppose the n-dimensional vector $\boldsymbol{Y} \sim \text{CM}\left(\mathbf{0}_{n,1}, \boldsymbol{V}, \boldsymbol{\alpha}_Y, \boldsymbol{\kappa}_Y\right)$, *where* $\boldsymbol{\alpha}_Y \equiv (\alpha_{Y,1}, \ldots, \alpha_{Y,n})'$,

Quantities Needed for Gibbs Sampling
$\boldsymbol{\mu}_\beta = (-\boldsymbol{\eta}'\boldsymbol{\Phi}' - \boldsymbol{\xi}', \mathbf{0}_{1,p})'$
$\boldsymbol{\mu}_\eta = (-\boldsymbol{\beta}'\mathbf{X}' - \boldsymbol{\xi}', \mathbf{0}_{1,r})'$
$\boldsymbol{\mu}_\xi = (-\boldsymbol{\beta}'\mathbf{X}' - \boldsymbol{\eta}'\boldsymbol{\Phi}', \mathbf{0}_{1,n})'$
$\boldsymbol{\mu}_{\gamma,i} = (\eta_i, \mathbf{0}_{1,i-1})'; \quad i = 2,\ldots,r$
$\mathbf{H}_\beta = (\mathbf{X}', \mathbf{V}_\beta^{-1\prime})'$
$\mathbf{H}_\eta = (\boldsymbol{\Psi}', \mathbf{V}_\eta^{-1\prime})'$
$\mathbf{H}_\xi = (\mathbf{I}_n, \mathbf{V}_\xi^{-1\prime})'$
$\mathbf{H}_{\gamma,i} = \left\{(\eta_1,\ldots,\eta_{i-1},)', \mathbf{C}_i'\right\}'; \quad i = 2,\ldots,r$
$\boldsymbol{\alpha}_\beta^* = (\mathbf{Z}' + \mathbf{d}_\beta', \boldsymbol{\alpha}_\beta - \mathbf{d}_\beta'\mathbf{X}\mathbf{V}_\beta)'$
$\boldsymbol{\alpha}_\eta^* = (\mathbf{Z}' + \mathbf{d}_\eta', \boldsymbol{\alpha}_\eta' - \mathbf{d}_\eta'\boldsymbol{\Phi}\mathbf{V}_\eta)'$
$\boldsymbol{\alpha}_\xi^* = (\mathbf{Z}' + \mathbf{d}_\xi', \boldsymbol{\alpha}_\xi' - \mathbf{d}_\xi'\mathbf{V}_\xi)'$
$\boldsymbol{\alpha}_{\gamma,i} = (\alpha_{\eta,i}, \boldsymbol{\alpha}_i')'; \quad i = 2,\ldots,r$
$\mathbf{d}_\beta \equiv \frac{\min(\boldsymbol{\alpha}_\beta)}{\min(\boldsymbol{\alpha}_\beta)+\max\left\{\text{abs}\left(\mathbf{J}_{n,1}'\mathbf{X}\mathbf{V}_\beta\right)\right\}}\mathbf{e}_0(\mathbf{Z}) - \frac{\min(\boldsymbol{\kappa}_\beta-\boldsymbol{\alpha}_\beta)}{\min(\boldsymbol{\kappa}_\beta-\boldsymbol{\alpha}_\beta)+\max\left\{\text{abs}\left(\mathbf{J}_{n,1}'\mathbf{X}\mathbf{V}_\beta\right)\right\}}\mathbf{e}_b(\mathbf{Z})$
$\mathbf{d}_\eta \equiv \frac{\min(\boldsymbol{\alpha}_\eta)}{\min(\boldsymbol{\alpha}_\eta)+\max\left\{\text{abs}\left(\mathbf{J}_{n,1}'\boldsymbol{\Phi}\mathbf{V}\right)\right\}}\mathbf{e}_0(\mathbf{Z}) - \frac{\min(\boldsymbol{\kappa}_\eta-\boldsymbol{\alpha}_\eta)}{\min(\boldsymbol{\kappa}_\eta-\boldsymbol{\alpha}_\eta)+\max\left\{\text{abs}\left(\mathbf{J}_{n,1}'\boldsymbol{\Phi}\mathbf{V}\right)\right\}}\mathbf{e}_b(\mathbf{Z})$
$\mathbf{d}_\xi \equiv \frac{\min(\boldsymbol{\alpha}_\xi)}{\min(\boldsymbol{\alpha}_\xi)+\max(\mathbf{J}_{n,1}'\mathbf{V}_\xi)}\mathbf{e}_0(\mathbf{Z}) - \frac{\min(\boldsymbol{\kappa}_\xi-\boldsymbol{\alpha}_\xi)}{\min(\boldsymbol{\kappa}_\xi-\boldsymbol{\alpha}_\xi)+\max(\mathbf{J}_{n,1}'\mathbf{V}_\xi)}\mathbf{e}_b(\mathbf{Z})$
$\boldsymbol{\kappa}_\beta^* = (\mathbf{J}_{1,p}, \kappa_\beta'\mathbf{J}_{p,1}')'$
$\boldsymbol{\kappa}_\eta^* = (\mathbf{J}_{1,r}, \boldsymbol{\kappa}_\eta')'$
$\boldsymbol{\kappa}_\xi^* = (\mathbf{J}_{1,n}, \kappa_\xi\mathbf{J}_{n,1}')'$
$\boldsymbol{\kappa}_{\gamma,i} = (\kappa_{\eta,i}, \boldsymbol{\kappa}_i')'; \quad i = 2,\ldots,r$

Table 3: A comprehensive list of matrices, vectors, and constants based on Theorems 2 and 5. Let $\mathbf{e}_0(\mathbf{Z}) \equiv (I(Z_1 = 0),\ldots,I(Z_n = 0))'$ if Z_i is bounded below by zero, and set $\mathbf{e}_0(\mathbf{Z}) \equiv \mathbf{0}_{n,1}$ otherwise. Let $\mathbf{e}_b(\mathbf{Z}) \equiv (I(Z_1 = b),\ldots,I(Z_n = b))'$ if Z_i is bounded above by b, and set $\mathbf{e}_b(\mathbf{Z}) \equiv \mathbf{0}_{n,1}$ otherwise.

$\boldsymbol{\kappa}_Y \equiv (\kappa_{Y,1}, \ldots, \kappa_{Y,n})'$, $\alpha_{Y,i}/\kappa_{Y,i} \in \mathscr{Y}$, *and* $\kappa_{Y,i} > 0$ *for* $i = 1, \ldots, n$. *Also assume the prior distribution in (16). Then, we have the following posterior distribution for* $\{\boldsymbol{v}_i\}$.

$$f(\boldsymbol{v}_2, \ldots, \boldsymbol{v}_n | \boldsymbol{Y}, \boldsymbol{\alpha}_Y, \boldsymbol{\kappa}_Y) = \prod_{i=2}^{n} f(\boldsymbol{v}_i | \boldsymbol{Y}, \boldsymbol{\alpha}_Y, \boldsymbol{\kappa}_Y) = \prod_{i=2}^{n} \mathrm{CM_c}(\boldsymbol{\mu}_{\gamma,i}, \boldsymbol{H}_{\gamma,i}, \boldsymbol{\alpha}_{\gamma,i}, \boldsymbol{\kappa}_{\gamma,i};\ \psi), \tag{17}$$

where $\boldsymbol{H}_{\gamma,i} \equiv \left(\boldsymbol{\Sigma}_i', \boldsymbol{C}_i'\right)'$, $\boldsymbol{\mu}_{\gamma,i} = (Y_i,\ \mathbf{0}_{1,i-1})'$, $\boldsymbol{\Sigma}_i' = \left(Y_j : j = 1, \ldots, i-1\right)'$, $\boldsymbol{\alpha}_{\gamma,i} = (\alpha_{Y,i}, \boldsymbol{\alpha}_i')'$, *and* $\boldsymbol{\kappa}_{\gamma,i} = (\kappa_{Y,i}, \boldsymbol{\kappa}_i')'$ *for* $i = 2, \ldots, n$

The CM prior distribution on the modified Cholesky decomposition of the precision matrix is similar to priors considered by Daniels and Pourahmadi (2002), Chen and Dunson (2003), and Pourahmadi et al. (2007) in the Gaussian setting. In fact, when $\psi = \psi_4$ the prior distribution in (16) reduces to the prior distributions used in Daniels and Pourahmadi (2002), Chen and Dunson (2003), and Pourahmadi et al. (2007). Thus, (16) constitutes a general non-Gaussian (natural exponential family) extension of such priors on modified Cholesky decompositions of precision and covariance matrices.

There are certainly other prior distributions for $\mathbf{V}^{-1}$ that may be more appropriate. For example, see Yang and Berger (1994) and Bradley et al. (2015c, for the spatial setting) for a Givens angle prior on covariance parameters. The Wishart and inverse Wishart are also common alternatives (e.g., see Gelman et al., 2013, for a standard reference). However, conjugacy may not always be present depending on the choice of CM distribution. Thus, in this article, we investigate the fully conjugate form of the LCM and specify the prior for $\mathbf{V}^{-1}$ as stated in (16).

2.5 Prior Distributions on DY Parameters

Bradley et al. (2015a) showed that the multivariate log-gamma distribution they proposed can be made arbitrarily close to a multivariate normal distribution by specifying the shape and scale parameters to be large. This essentially allows one to use a LGP specification with a Poisson data

model, and *also use* the conjugacy that arises from the MLG distribution when $j = k = 3$ in (14). This important property of the MLG distribution can be extended to the more general CM distribution.

Theorem 5: Suppose that $\psi \neq \psi_4$, and denote the first and second derivatives with ψ' and ψ'', $0 < \psi' < \infty$, and $0 < \psi'' < \infty$. Let the n-dimensional random vector $\boldsymbol{Y} \sim \mathrm{CM}(\boldsymbol{\mu}, \left(\psi''(0)/\psi'(0)\right)^{1/2} \alpha^{1/2}\boldsymbol{V}, \alpha\boldsymbol{J}_{n,1}, \frac{\alpha}{\psi'(0)}\boldsymbol{J}_{n,1};\ \psi)$. Then $\boldsymbol{Y}$ converges in distribution to a multivariate normal random vector with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{V}\boldsymbol{V}'$ as α approaches infinity.

The restriction of $\psi \neq \psi_4$ is sensible, since $\psi = \psi_4$ yields a CM exactly equal to a multivariate normal distribution. Also, Theorem 5 does not hold for the multivariate negative-inverse-gamma distribution, since $\psi_1'(0) = -\infty$.

The "best" DY parameters, for the multivariate logit-beta distribution and the MLG distribution, might not lead to something that looks Gaussian. That is, we should be able to learn whether or not the multivariate normal distribution is appropriate for latent processes of binomial and Poisson data by observing whether or not posterior replicates of the DY parameters (i.e., α and κ) are large (which would invoke Theorem 5). Hence, from this point-of-view, it is very important that we place prior distributions on the DY parameters.

Following the theme of the previous sections, we define conjugate priors for the DY parameters by defining a distribution with an exponential term and an exponential to the negative unit log partition function. That is, consider

$$f\left(\alpha_\xi, \kappa_\xi | \boldsymbol{\gamma}_\xi \rho_\xi\right) \propto \exp\left[\gamma_1 \alpha + \gamma_2 \kappa - \rho_\xi \log\left\{\frac{1}{K(\alpha, \kappa)}\right\}\right], \tag{18}$$

where γ_1, γ_2, and ρ are hyperparameters. The parameter space for γ_1, γ_2, and ρ that ensures that (18) is proper (i.e., can be normalized to define a probability measure) is an immediate consequence

Unit Log Partition Function	**Form of the Prior Distribution on α and κ (i.e., $f(\alpha,\kappa\|\gamma_1,\gamma_2,\rho)$)**	**Suggested Hyperparameters**	**Special Case of the Prior Distribution**
$\psi_1(Y)=\log\left(-\frac{1}{Y}\right)$	$\exp\{\gamma_1\alpha+\gamma_2\kappa-\rho\log(\Gamma(\kappa+1))-\rho(\kappa+1)\log(\alpha)\}$ $=\frac{1}{\Gamma(\kappa+1)^\rho}\left(\alpha^{-\rho}\exp(\gamma_2)\right)^{\kappa+1}\exp(\gamma_1\alpha)$	$\gamma_1=-1000$ $\gamma_2=1000$ $\rho=10^{-15}$	If κ is integer-valued then the conditional distribution of $\kappa\|\alpha$ is Conway-Maxwell-Poisson with parameters $\alpha^\rho\exp(\gamma_2)$ and ρ, and the conditional distribution of $\alpha\|\kappa$ is Gamma$((\kappa+1)\rho+1,-1/\gamma_1)$ provided that $\gamma_2\in\mathbb{R}$, γ_1 is negative, and $\rho\geq 0$.
$\psi_2(Y)=\log(1+\exp(Y))$	$\exp[\gamma_1\alpha+\gamma_2\kappa+\rho\log\{\Gamma(\kappa)\}-\rho\log\{\Gamma(\alpha)\}-\rho\log\{\Gamma(\kappa-\alpha)\}]$ $=\left(\frac{\Gamma(\kappa)}{\Gamma(\alpha)\Gamma(\kappa-\alpha)}\right)^\rho\exp(\gamma_1)^\alpha\exp(\gamma_2)^\kappa$	$\gamma_1=0$ $\gamma_2=-1000$ $\rho=1$	Let $\rho=1$, $\gamma_1\in\mathbb{R}$, and $\gamma_2<0$. If α and κ are integer-valued, then the conditional distribution of $(\alpha-1)\|\kappa$ is binomial with κ number of Bernoulli trials, and probability of success $\exp(\gamma_1)/(1+\exp(\gamma_1))$. Also, $(\kappa-\alpha-1)\|\alpha$ follows a negative binomial distribution with $\alpha+1$ number of successful Bernoulli trials, and probability of success $\exp(\gamma_2)$.
$\psi_3(Y)=\exp(Y)$	$\exp\{\gamma_1\alpha+\gamma_2\kappa-\rho\log(\Gamma(\alpha))-\rho(\alpha)\log(\kappa)\}=\frac{1}{\Gamma(\alpha)^\rho}\left(\kappa^{-\rho}\exp(\gamma_1)\right)^\alpha\exp(\gamma_2\kappa)$	$\gamma_1=1000$ $\gamma_2=-10^{-15}$ $\rho=1$	If α is integer-valued then the conditional distribution of $(\alpha-1)\|\kappa$ is Conway-Maxwell-Poisson with parameters $\kappa^\rho\exp(\gamma_1)$ and ρ, and the conditional distribution of $\kappa\|\alpha$ is Gamma$(\alpha\rho+1,-1/\gamma_2)$ provided that $\gamma_1\in\mathbb{R}$, γ_2 is negative, and $\rho\geq 0$.
$\psi_4(Y)=Y^2$	$\exp\left(\gamma_1\alpha+\gamma_2\kappa+\frac{\rho}{2}\log\kappa-\frac{\alpha^2}{4\kappa}\right)=\kappa^{\rho/2+1-1}\exp(\gamma_2\kappa)\exp(-\frac{(\alpha-2\kappa\gamma_1)^2}{4\kappa})$	Set $\alpha=0$ $\gamma_1=0$ $\gamma_2=-\frac{1}{2}$ $\rho=2$	We have that κ is distributed as Gamma$(\rho/2+1,-1/\gamma_2)$ and is independent of α, which is distributed as normal with mean $2\kappa\gamma_1$ and variance 2κ. The suggested hyperparameters result in an inverse-gamma prior distribution on the variance of a normal random variable with shape 2 and scale 1, which yields mean 1 and variance infinity.

Table 4: Special Cases: We list the form of the the prior distribution in (18) by ψ_j for $j=1,\ldots,4$. The first column has the unit log partition function, the second column has the form of the prior distributions (up to a proportionality constant), the third column gives suggested hyperparameters, and the fourth column gives special cases of the conditional distributions $\alpha|\kappa$ and $\kappa|\alpha$.

of a result from Diaconis and Ylvisaker (1979). In particular, from Theorem 1 of Diaconis and Ylvisaker (1979), the distribution in (18) is proper provided that $\mathscr{Y}$ is a nonempty real-valued open set, the range of ψ is a nonempty real-valued open set, $\gamma_1/\rho \in \mathscr{Y}$, $\gamma_2/\rho \in \mathscr{Y}_\psi$, and $\rho > 0$, where $\mathscr{Y}_\psi \equiv \{M : M = -\psi(Y), Y \in \mathscr{Y}\}$. For the CM distribution associated with ψ_1 we see that $\mathscr{Y} = \{Y : Y < 0\}$ and the range of ψ is $\mathbb{R}$; thus, for this setting $\gamma_1 < 0$, $\gamma_2 \in \mathbb{R}$, and $\rho > 0$ results in a proper prior in (18). For the CM distributions associated with ψ_2, ψ_3, and ψ_4 we have that $\mathscr{Y} = \mathbb{R}$ and ψ is a strictly positive; thus, for this setting $\gamma_1 \in \mathbb{R}$, $\gamma_2 < 0$, and $\rho > 0$ ensures propriety of (18).

There are many interesting special cases of the prior distribution in (18). For example, when α is integer-valued and $\psi = \psi_3$ then the prior in (18) has a relationship with the Conway-Maxwell-Poisson distribution (Conway and Maxwell, 1962) and the gamma distribution. These special cases (listed in Table 4 for ease of exposition) are particularly useful because they give rise to interpretations of the hyperparameters. In particular, for $\psi = \psi_1$, we have that ρ can be interpreted as a dispersion parameter (in relation to the dispersion parameter of a Conway-Maxwell-Poisson distribution), γ_1 can be interpreted as a location parameter, and γ_2 can be interpreted as a scale parameter. When $\psi = \psi_2$ we have that γ_1 and γ_2 can be interpreted as functions of a proportion (i.e., the inverse logit or log of a proportion). For $\psi = \psi_3$, we have that ρ can be interpreted as a dispersion parameter, γ_2 can be interpreted as a location parameter, and γ_1 can be interpreted as a scale parameter. Finally, when $\psi = \psi_4$ we have that γ_1 is interpreted as a location parameter, ρ represents a shape parameter, and γ_2 represents a scale parameter.

The most familiar special case occurs when $\psi = \psi_4$ (i.e., a normal data model) and $\alpha = 0$. Namely, (18) reduces to independent gamma prior distributions on κ with shape parameter $\rho/2+1$, and scale parameter $-\gamma_2$. When recognizing that κ is equal to one-half the unknown variance of a normal random variable (See Table 1), we see that the conjugate prior distribution implies an inverse gamma distribution for the variance parameter, which is a common choice of a prior distribution on the variance parameter for normally distributed data (Gelman, 2006).

Now, consider $\mathbf{q} \sim \text{CM}(\mathbf{0}_{r,1}, \mathbf{V}_q, (\alpha_1, \ldots, \alpha_r)', (\kappa_1, \ldots, \kappa_r)')$, $\mathbf{V}_q \in \mathbb{R}^r \times \mathbb{R}^r$, $\mathbf{Y} \sim \text{CM}(\mathbf{0}_{n,1}, \mathbf{I}_n, \alpha\mathbf{J}_{n,1}, \kappa\mathbf{J}_{n,1})$, (α_i, κ_i) are independently distributed according to (18) for $i = 1, \ldots, r$, and (α, κ) are distributed according to (18). Then full conditional distributions associated with $\{\alpha_i, \kappa_i\}$ and (α, κ) are immediate:

$$
\begin{aligned}
& f(\alpha_i, \kappa_i | \cdot) \\
& \propto \exp\left[(\gamma_1 + \mathbf{j}_i' \mathbf{V}_q^{-1} \mathbf{q})\alpha_i + (\gamma_2 - \mathbf{j}_i' \boldsymbol{\psi}(\mathbf{V}_q^{-1} \mathbf{q}))\kappa_i - (\rho + 1)\log\left\{\frac{1}{K(\alpha_i, \kappa_i)}\right\}\right]; \; i = 1, \ldots, r, \\
& f(\alpha, \kappa | \cdot) \\
& \propto \exp\left[(\gamma_1 + \mathbf{J}_{1,n}\mathbf{Y})\alpha + (\gamma_2 - \mathbf{J}_{1,n}\psi(\mathbf{Y}))\kappa - (\rho + 1)\log\left\{\frac{1}{K(\alpha, \kappa)}\right\}\right], \qquad (19)
\end{aligned}
$$

where $\mathbf{j}_i$ is the i-th column of $\mathbf{I}_n$. These full conditional distributions are straightforward to simulate from using a slice sampler (Neal, 2003).

2.6 Model Summary

Using (14), (15), and (19), we can write the joint distribution of the data, random processes, covariances, and DY parameters as the product of the following conditional and marginal distributions:

$$
\begin{aligned}
&\text{Data Model}: \; Z_i|\boldsymbol{\beta},\boldsymbol{\eta},\xi_i \overset{\text{ind}}{\sim} \text{EF}\left(\mathbf{x}_i'\boldsymbol{\beta}+\boldsymbol{\phi}_i'\boldsymbol{\eta}+\xi_i;\; \psi_j\right); \quad i=1,\ldots,n, j=1,2,3,4\\
&\text{Process Model 1}: \; \boldsymbol{\eta}|\mathbf{V},\alpha_\eta,\kappa_\eta \sim \text{CM}\left(\mathbf{0}_{r,1},\mathbf{V},\boldsymbol{\alpha}_\eta,\boldsymbol{\kappa}_\eta;\; \psi_k\right);\\
&\text{Process Model 2}: \; \boldsymbol{\xi}|\alpha_\xi,\kappa_\xi \sim \text{CM}\left(\mathbf{0}_{n,1},\mathbf{I}_n,\alpha_\xi\mathbf{J}_{n,1},\kappa_\xi\mathbf{J}_{n,1};\; \psi_k\right);\\
&\text{Parameter Model 1}: \; \boldsymbol{\beta} \sim \text{CM}\left(\mathbf{0}_{p,1},\sigma_\beta\mathbf{I}_p,\alpha_\beta\mathbf{J}_{p,1},\kappa_\beta\mathbf{J}_{p,1};\; \psi_k\right),\\
&\text{Parameter Model 2}: \; \mathbf{v}_i \overset{\text{ind}}{\sim} \text{CM}(\mathbf{0},\sigma_v\mathbf{I}_{i-1},\alpha_v\mathbf{J}_{i-1,1},\kappa_v\mathbf{J}_{i-1,1};\; \psi_k);\; i=2,\ldots,r, k=1,2,3,4,\\
&\text{Parameter Model 3}: \; f(\alpha_{\eta,m},\kappa_{\eta,m}|\gamma_{\eta,1},\gamma_{\eta,2},\rho_\eta)\\
&\qquad\qquad \propto \exp\left[\gamma_{\eta,1}\alpha_{\eta,m}+\gamma_{\eta,2}\kappa_{\eta,m}-\rho_\eta\log\left\{\frac{1}{K\left(\alpha_{\eta,m},\kappa_{\eta,m}\right)}\right\}\right];\; m=1,\ldots,r\\
&\text{Parameter Model 4}: \; f(\alpha_\xi,\kappa_\xi|\gamma_{\xi,1},\gamma_{\xi,2},\rho_\xi) \propto \exp\left[\gamma_{\xi,1}\alpha_\xi+\gamma_{\xi,2}\kappa_\xi-\rho_\xi\log\left\{\frac{1}{K\left(\alpha_\xi,\kappa_\xi\right)}\right\}\right],
\end{aligned}
\tag{20}
$$

where notice that "Process Model 1" and "Parameter Model 1" in (20) are less general than what is presented in (14). In our experience, we found that it is only necessary to specify $\boldsymbol{\eta}$ to have cross-covariances and heterogeneous specifications of the DY parameters. However, this may not be the case for every problem, and one can easily substitute the more general specification of "Process Model 1" and "Parameter Model 1" from (14) into (20).

To use this model in practice we need to specify the hyperparameters, α_β, κ_β, α_v, κ_v, σ_v, σ_β, $\gamma_{\eta,1}$, $\gamma_{\eta,2}$, $\gamma_{\xi,1}$, $\gamma_{\xi,2}$, ρ_η, and ρ_ξ. In general, we suggest specifying the hyperparameters so that the corresponding parameter models are flat or "vague." In particular, we choose $\alpha_\beta=\kappa_\beta=\alpha_v=\kappa_v=500$, and $\sigma_\beta=\sigma_v=100$, and have found that our results are robust to these specifications. Additionally, the suggested choices of hyperparameters associated with $\{\alpha_{\eta,i},\kappa_{\eta,i}\}$, α_ξ, and κ_ξ are outlined in Table 4 by the appropriate CM distribution. These choices were made so that the

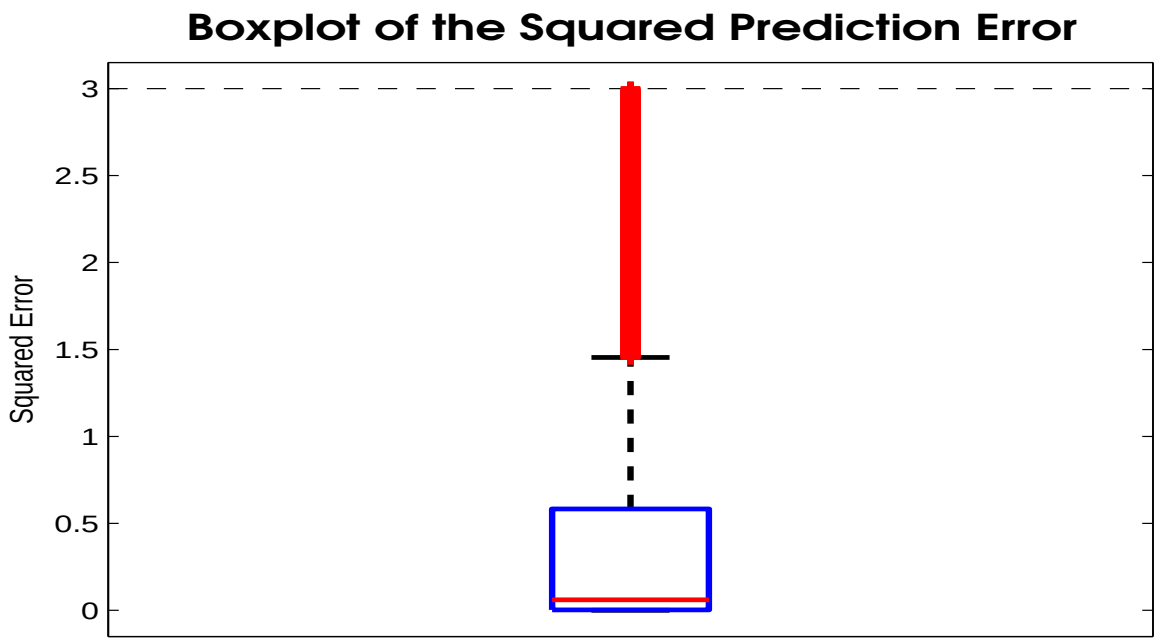

Figure 1: Boxplot of prediction error: We simulate $n = 1,000,000$ values as described in Section 3, and fit an LCM with the binomial data model and $j = k = 3$. The boxplot is over the 1,000,000 values of $(tp_1 - t\widehat{p}_1)^2, \ldots, (tp_n - t\widehat{p}_n)^2$ for $n = 1,000,000$. Outliers above three are omitted for the purposes of visualization.

special cases listed in the fourth column of Table 4 are "flat" or "vague." In the Supplemental Appendix, we provide step-by-step instructions on implementing a Gibbs sampler for the model in (20).

3 Simulation Studies

We provide comparisons of predictions using the LGP (i.e., the LCM in (20) with $j = 2$ and $k = 4$) versus predictions based on a LCM in (20) with $j = k = 2$. As discussed in Section 1, the LGP is the standard approach for Bayesian analysis of dependent data, and thus, the results in this section are meant to provide one comparison of the LCM to the current state-of-the-art. We find it important to emphasize that if the LGP is more appropriate than the LCM, our model will be able to identify

this because of Theorem 5; that is, if the posterior replicates of the DY parameters are large then Theorem 5 suggests that the latent processes are approximately Gaussian.

The $n \times p$ matrix $\mathbf{X} \equiv (\mathbf{x}_1, \ldots, \mathbf{x}_p)'$, the $r \times r$ matrix $\boldsymbol{\Phi} \equiv (\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_r)'$, and the $r \times r$ lower unit triangle matrix $\mathbf{V}^{-1}$ are randomly generated with $p = 3$ and $r = 10$. Specifically, each non-zero or one element of the $n \times 3$ matrix $\mathbf{X}$, the $n \times 10$ matrix $\boldsymbol{\Psi}$, and the 10×10 matrix $\mathbf{V}^{-1}$ are selected from a standard normal distribution. The elements of the fixed and random effects $\boldsymbol{\beta}$, $\boldsymbol{\eta}$, and $\boldsymbol{\xi}$ are randomly selected from a standard normal distribution as well. Then, we define $p_i = \frac{\exp(\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\phi}_i'\boldsymbol{\eta} + \xi_i)}{1+\exp(\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\phi}_i'\boldsymbol{\eta} + \xi_i)}$ for $i = 1, \ldots, n$. We consider two different data models in this section. In particular, we consider observations Z_i generated from a binomial distribution with sample size $t = 10$ and probability of success p_i for $i = 1, \ldots, n$. (For example, consider n households, that consists of t individuals, and let p_i represent the probability that an individual is female. Then Z_i would represent the number of women living in household i.) Similarly, consider observations Z_i generated from a negative binomial distribution with $t = 10$ successes and probability of success p_i for $i = 1, \ldots, n$.

Using the Gibbs sampler outlined in the Supplemental Appendix we implement the LCM in (20) with $j = k = 2$ and use the appropriate data model (i.e., binomial or negative binomial). To fit the LGP, we use the Matlab (Version 8.0) function *slicesample* on the joint pdf stated in (20) with $j = 2$ and $k = 4$ and the appropriate data model. The choice of hyperparameters discussed in Section 2.6 were used. For each i, denote the posterior mean of p_i with $\widehat{p}_i$, the posterior mean of $\text{odds}_i \equiv p_i/(1-p_i)$ with $\widehat{\text{odds}}_i$, and define the total squared prediction errors to be

$$\sum_i (t p_i - t\widehat{p}_i)^2$$
$$\sum_i \left(\text{odds}_i - \widehat{\text{odds}}_i\right)^2,$$

used for the binomial and negative binomial settings, respectively. We used a burn-in of 10,000 and generate $B = 20,000$ posterior replications for both data models that are considered.

We have not yet specified the choice of n, and we begin with a demonstration from this simulation model with a single replication of $\{Z_1,\ldots,Z_n\}$ with $n = 1,000,000$. This is a massive data example, which is becoming a standard in modern statistics. In Figure 1, we provide a boxplot of the values $(tp_1 - t\widehat{p}_1)^2,\ldots,(tp_n - t\widehat{p}_n)^2$ for $n = 1,000,000$. In general the predictions are very close (in squared error) to the true values of $tp_1,\ldots,tp_n$, since around 75% of the data have a squared error of less than one. The Gibbs sampler took approximately 7 hours to execute, where all of our computations (in both Sections 4 and 5) were performed in Matlab (Version 8.0) on a dual 10 core 2.8 GHz Intel Xeon E5-2680 v2 processor with 256 GB of RAM. Visually, Figure 1 suggests accurate predictions; however, a very basic way to evaluate the predictions is to compare to the total prediction error when using the original dataset as a predictor (i.e., compute $\sum_i (tp_i - t\widehat{p}_i)^2$ and $\sum_i (tp_i - Z_i)^2$). For this example, $\sum_i (tp_i - t\widehat{p}_i)^2 = 7.02 \times 10^5$ and $\sum_i (tp_i - Z_i)^2 = 9.32 \times 10^5$. Thus, this simulation suggests that the LCM model yields precise predictions, and is computationally feasible for extremely large datasets. Note that the high predictive performance of the LCM model occurs in a setting where we do not generate the truth from a multivariate logit-beta distribution.

The advantage of the LCM model immediately becomes apparent when attempting to compare to predictions from an LGP model. For $n = 1,000,000$ the LGP model becomes extremely difficult to fit. Thus, to compare the LCM to the LGP we set $n = 30$ (a small data problem). Certainly, the LGP can be implemented on a dataset much larger than $n = 30$, but a dataset of this size facilitates comparison of the LCM model to the LGP model over multiple replicates of the dataset $\{Z_1,\ldots,Z_n\}$. Now, we simulate the dataset $\{Z_1,\ldots,Z_n\}$ fifty times from the binomial distribution, and simulate another fifty independent replications of of $\{Z_1,\ldots,Z_n\}$ from the negative binomial distribution. In Figure 2(a,b) we plot the ratio of total squared prediction error of the LCM model and the total squared prediction error of the LGP model over the fifty independent replicates of $\{Z_i\}$ by the corresponding data model. (For the negative binomial setting we plot the log of the ratio of the total square prediction errors for the purpose of visualization.) In Figure 2(a) a log ratio greater than zero indicates that the LCM has smaller total square prediction error. Likewise, in Figure 2

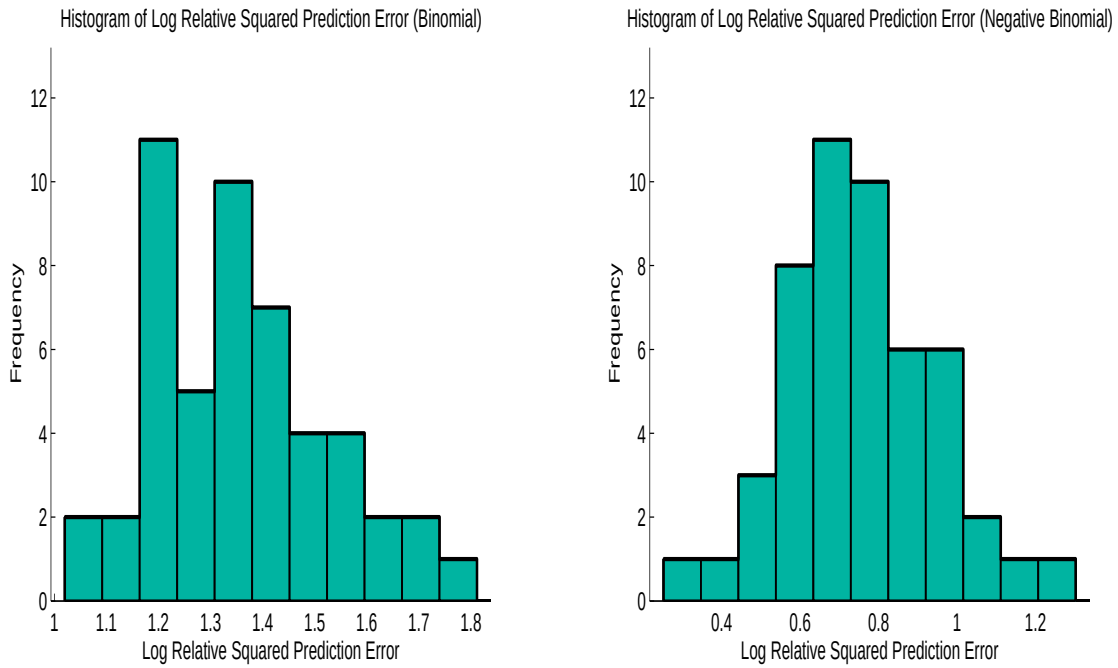
Histogram of Log Relative Squared Prediction Error (Binomial)
Frequency
Log Relative Squared Prediction Error
Histogram of Log Relative Squared Prediction Error (Negative Binomial)
Frequency
Log Relative Squared Prediction Error

Figure 2: Histogram over log relative total square prediction error: For each of the fifty realizations of $\{Z_1, \ldots, Z_{30}\}$ from a (negative) binomial distribution, we produce $(\{\widehat{\text{odds}}_i\})$ $\{\widehat{p}_i\}$ using the appropriate LCM model, produce $(\{\widehat{\text{odds}}_i\})$ $\{\widehat{p}_i\}$ using the LGP model, and compute the log relative total squared prediction error. The log relative total squared prediction error for the (negative) binomial setting is the log of the total squared prediction error of $(\{\widehat{\text{odds}}_i\})$ $\{t\widehat{p}_i\}$ from the LGP model minus the log total squared prediction error of $(\{\widehat{\text{odds}}_i\})$ $\{t\widehat{p}_i\}$ computed using the LCM model. The histogram in the (right) left panel is over the 50 independent replicates from the (negative) binomial distribution.

(b), a log ratio greater than zero indicates that the LCM has smaller total squared prediction error. Here, we see that the ratio of the total squared prediction errors are consistently larger than zero, and hence, the LCM clearly outperforms the LGP for this simulation design for both the binomial and negative binomial settings. Thus, not only does the LCM lead to computational advantages over the LGP for this example, there are also clear gains in predictive performance.

4 An Application to Count-Valued ACS Public-Use Micro-Data

The US Census Bureau has replaced the decennial census long-form with the American Community Survey (ACS), which is an ongoing survey that collects an enormous amount of information on US demographics. (To date there are over 64,000 variables published through the ACS.) The estimates published from the ACS have a unique multi-year structure. Specifically, the ACS produces 1-year and 5-year period estimates of US demographics, where 1-year period estimates are summaries (e.g., median income of a particular county) made available over populations over 65,000 and 5-year period estimates are made available for all published geographies (e.g., see Torrieri, 2007, for more information).

A difficulty with using ACS period estimates published over pre-defined geographies is that it is difficult to infer fine-level (i.e., household) information. As a result, the ACS provides a public−use micro-sample (PUMS) over public−use micro-areas (PUMAs). PUMS consists of individual and household information within each PUMA, where the location of the household within the PUMA is not released to the public. In this section, we focus on household level PUMS found within one particular PUMA; namely the PUMA that covers the metropolitan area of Tallahassee Florida (labeled as PUMA number 00701). Consider 2005-2009 PUMS estimates of the number of individuals living in a household contained within PUMA 00701. This is a fairly large dataset (for multivariate statistics) consisting of 4,537 observations. An important inferential goal, besides giving an illustration of the LCM, is to accurately predict vacant households (i.e., predict zero people living in a household). Vacant households exhaust resources for those conducting surveys, and is of practical interest to the US Census Bureau (see, `http://www.census.gov/en.html`).

We would expect the number of individuals living in a household to be spatially correlated, since certain neighborhoods within Tallahassee are known to be more attractive for those with a family, and hence, have more people living within a household in these neighborhoods. However, the spatial correlation can not be leveraged, since the location *within* PUMAs are not publicly

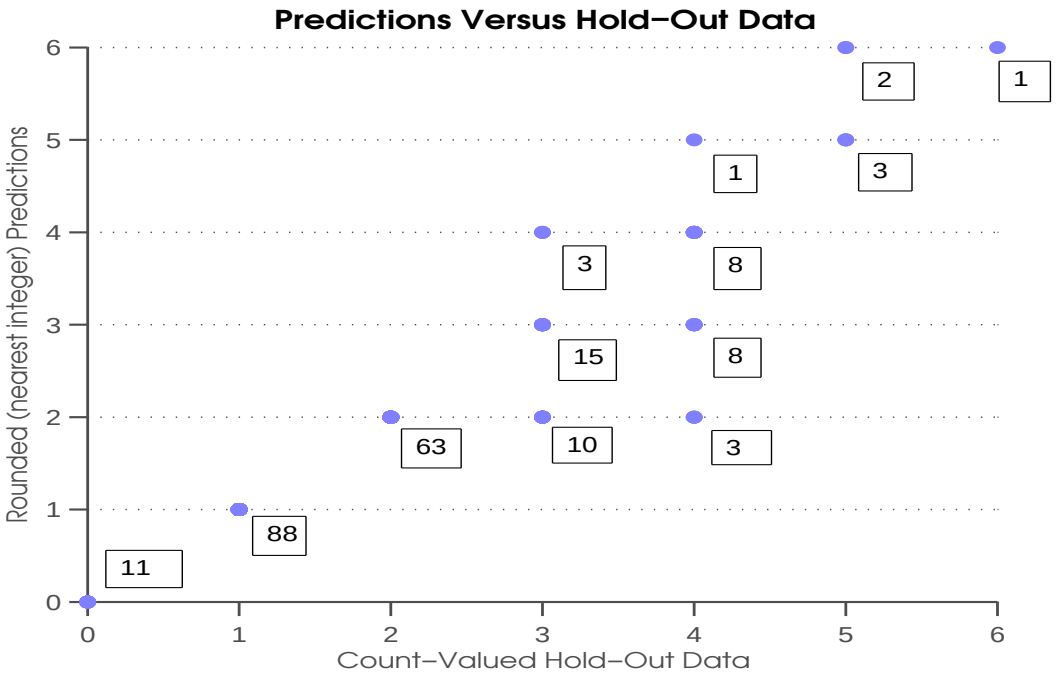

Figure 3: A scatterplot of the predicted mean for the randomly selected hold-out dataset with 216 observations. These predictions are rounded to the nearest integer, since the hold-out dataset is known to be integer-valued. Each point has a text-box to it's immediate bottom-right, except the point at the origin, where the text-box is to it's immediate top-right. The numbers in this textbox indicate the number of instances (out of the 216 hold-out values) where this point occurred.

available. Consequently, we model the dependencies within the PUMS using a generic multivariate distribution; namely, the CM distribution. In particular, we assume that the data follows a LCM with $j = k = 3$; that is, a Poisson data model with multivariate log-gamma process models. There are a large number of potential covariates (there are 358 in total) including fuel cost of the household, number of bedrooms in the household, and lot size, among others. For illustration, we picked a small subset of covariates using least angle regression (Efron et al., 2004), which lead to $p = 8$ covariates. Namely, we use the Moran's I basis functions (Griffith, 2000, 2002, 2004; Hughes and Haran, 2013; Bradley et al., 2015b), where ϕ_i is set equal to the i-th eigenvector of $\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}$. We sequentially considered different values of r (starting with $r = 0, 1, 2, \ldots$), and computed the deviance information criterion (DIC; Spiegelhalter et al., 2002) associated with the

LCM, and found that $r = 5$ is a reasonable choice. The hyperparameters discussed in Section 2.6 were used. All of our computations (in both Sections 4 and 5) were performed in Matlab (Version 8.0) on a dual 10 core 2.8 GHz Intel Xeon E5-2680 v2 processor with 256 GB of RAM. Convergence of the MCMC algorithm was assessed visually using trace plots, and no lack of convergence was detected.

To assess the quality of the predictions we randomly selected 216 observations (roughly 5% of the data). Using the remaining data we produce estimates of the mean number of individuals living in a household for the 216 observations. In Figure 3 we plot a scatterplot of the hold-out dataset and the corresponding predictions that were based on the remaining $4,321$ observations. Here, we see that a majority of the rounded (to the nearest integer) predictions are *exactly equal* to the hold-out data. In fact, $189/216 \approx 88\%$ of the 216 hold-out dataset are *exactly* equal to the corresponding rounded predicted value, and the remaining 12% are within one count of the corresponding hold-out data value. Furthermore, we are able to very accurately predict an empty household, which may have implications for sampling done by the US Census Bureau.

5 Discussion

We have introduced methodology for jointly modeling dependent non-Gaussian data within the Bayesian framework. This methodology is rooted in the development of new distribution theory for dependent data that makes Bayesian inference possible to implement using a Gibbs sampler; hence, computationally intensive and ad hoc approaches needed for tuning and specifying proposal distributions are not needed. Specifically, we propose a multivariate version of the prior distributions introduced by Diaconis and Ylvisaker (1979). Furthermore, the prior distributions similar to those used by Daniels and Pourahmadi (2002), Chen and Dunson (2003), and Pourahmadi et al. (2007) are adapted to the non-Gaussian setting.

Several theoretical results were required to derive this *conjugate multivariate distribution*

(CM), and to develop its use for Bayesian inference of dependent data from the natural exponential family. The later is facilitated through the introduction of the latent CM (LCM) model. In particular, we show that full conditional distributions are of the same form of a conditional distribution of a CM random vector, and provide a way to simulate from this conditional distribution. Relationships between the LCM and the LGP also provide motivation for the use of the LCM. In particular, the latent Gaussian process (LGP) is a special case of the LCM. Furthermore, many types of LCMs can be well approximated by a LGP, by specifying certain parameters of a LCM to be "large." This result shows that the LCM is not only computationally easier to implement, but is also more flexible than a LGPs.

Empirical exploration of the Poisson, binomial, and negative binomial special cases were performed through simulations studies and through an analysis of period estimates made available by the US Census Bureau's American Community Survey (ACS). These examples indicate very small out-of-sample error when using LCM for prediction, and show gains in predictive performance over the LGP. Additionally, the LCM model is applicable for extremely large datasets (in the simulation study we implemented the LCM on a simulated dataset of size 1,000,000), while the LGP can be extremely difficult to implement in high dimensions. In the real data analysis section we predict the number of individuals within a household over the US city of Tallahassee Florida, and obtain very precise estimates (in terms of hold-out error). The predictions were very accurate even though $11/216 \approx 6\%$ of the hold-out dataset consisted of zero counts, which is known to cause difficulties in an LGP (Lambert, 2006).

Although there are many settings where the LCM improves both precision and computation, there are settings where it would not be feasible to implement the LCM. In particular, we consider one choice of ψ that results in a case where $\mathscr{M}_p^{n+p}$, $\mathscr{M}_r^{n+r}$, and $\mathscr{M}_n^{2n}$ does not guarantee that $\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y}$ for each i; namely ψ_1, which is the unit log partition function of a gamma data model. In this case, the full conditional distributions in (15) are *truncated* CM_c distributions. Thus, in this setting the LCM is most easily implemented by doing a Gibbs sampler with component-

wise updates due to the truncated support of the natural parameter. This is computationally less efficient than simply transforming the gamma data to the log scale and fitting an LGP, which can give precise predictions.

As discussed in the Introduction, a general modeling framework for dependent data that can model non-Gaussian (natural exponential family) data as easily as Gaussian data, has important implications for applied statistics. Nevertheless, there are also many opportunities for new methodological results are exciting, since a special case of our framework (i.e., the LGP) has been the central methodological tool used in the dependent data literature. In particular, we are interested in developing the LCM model within "more specific" dependent data settings such as spatial, spatio-temporal, and multivariate spatio-temporal arenas.

Acknowledgments

We would like to thank Drs. Matthew Simpson and Erin Schliep at the University of Missouri for helpful discussions. This research was partially supported by the U.S. National Science Foundation (NSF) and the U.S. Census Bureau under NSF grant SES-1132031, funded through the NSF-Census Research Network (NCRN) program.

Supplemental Appendix: Bayesian Hierarchical Models with Conjugate Full-Conditional Distributions for Dependent Data from the Natural Exponential Family

Introduction

In this Supplemental Appendix, we provide the proofs of the proposition and theorems stated in the main text (see Appendix A). Additionally, we provide step-by-step instructions on the Gibbs sampler for the LCM model (see Appendix B).

Appendix A: Proofs

In this appendix we provide proofs for the technical results stated in the paper.

Proof of Proposition 1(i): From (2) of the main text we see that the distribution of the random vector $\mathbf{w}$ in (7) is given by,

$$\left(\prod_{i=1}^{n} K(\alpha_i, \kappa_i)\right) \exp\left\{\boldsymbol{\alpha}'\mathbf{w} - \boldsymbol{\kappa}'\psi(\mathbf{w})\right\}; \quad \mathbf{w} \in \mathbb{R}^n.$$

The inverse of the transform of (7) is given by $\mathbf{w} = \mathbf{V}^{-1}(\mathbf{Y} - \boldsymbol{\mu})$, and the Jacobian is given by $|\det(\mathbf{V}^{-1})|$. Then, by a change-of-variables (e.g., see Casella and Berger, 2002), we have that the pdf of $\mathbf{Y}$ is given by,

$$\det(\mathbf{V}^{-1}) \left(\prod_{i=1}^{n} K(\alpha_i, \kappa_i)\right) \exp\left[\boldsymbol{\alpha}'\mathbf{V}^{-1}(\mathbf{Y} - \boldsymbol{\mu}) - \boldsymbol{\kappa}'\psi\left\{\mathbf{V}^{-1}(\mathbf{Y} - \boldsymbol{\mu})\right\}\right]; \ \mathbf{Y} \in \mathscr{M}^n.$$

This completes the proof of Proposition 1(i).

Proof of Theorem 1: It follows from Proposition 1(i) that the conditional distribution is given by

$$\begin{aligned} f(\mathbf{Y}_1|\mathbf{Y}_2,\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) &= \frac{[f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})]_{\mathbf{Y}_2=\mathbf{d}}}{[\int f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})d\mathbf{Y}_1]_{\mathbf{Y}_2=\mathbf{d}}}, \\ &\propto \exp\left[\boldsymbol{\alpha}'(\mathbf{H}\ \ \mathbf{B})\begin{pmatrix}\mathbf{Y}_1\\ \mathbf{d}\end{pmatrix} - \boldsymbol{\kappa}'\psi\left\{(\mathbf{H}\ \ \mathbf{B})\begin{pmatrix}\mathbf{Y}_1\\ \mathbf{d}\end{pmatrix} - \mathbf{V}^{-1}\boldsymbol{\mu}\right\}\right], \\ &\propto \exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi\left(\mathbf{H}\mathbf{Y}_1+\mathbf{B}\mathbf{d}-\mathbf{V}^{-1}\boldsymbol{\mu}\right)\right\}, \\ &= \exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi(\mathbf{H}\mathbf{Y}_1-\boldsymbol{\mu}^*)\right\};\ \mathbf{Y}_1\in\mathbb{R}^n, \end{aligned}$$

which proves the result. The normalizing constant can be found using a change of variables

$$M = \frac{\det(\mathbf{V}^{-1})\left\{\prod_{i=1}^n K(\alpha_i,\kappa_i)\right\}\exp\left(\boldsymbol{\alpha}'\mathbf{B}\mathbf{d}-\boldsymbol{\alpha}'\mathbf{V}^{-1}\boldsymbol{\mu}\right)}{[\int f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})d\mathbf{Y}_1]_{\mathbf{Y}_2=\mathbf{d}}} \tag{A.1}$$

Although we do not find the expression of the integral $[\int f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})d\mathbf{Y}_1]_{\mathbf{Y}_2=\mathbf{d}}$, and consequently M, we know that M is non-zero and finite. To see this, let $\mathscr{N}_1=\{\mathbf{Y}_2:[\int f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})d\mathbf{Y}_1]_{\mathbf{Y}_2}=0\}$; then, by the definition of the CM distribution for $\mathbf{Y}\in\mathscr{M}^n$ and $\mathbf{Y}_2\in\mathscr{N}_1$

$$f\left(\begin{bmatrix}\mathbf{Y}_1\\ \mathbf{Y}_2\end{bmatrix}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}\right) > 0.$$

Taking the integral with respect to $\mathbf{Y}_1$ on both sides of the inequality gives $0>0$, which is a false statement. Thus, we have that $[\int f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})d\mathbf{Y}_1]_{\mathbf{Y}_2=\mathbf{d}}$ is non-zero, and hence, M is finite. Similarly, let $\mathscr{N}_2=\{\mathbf{Y}_2:[\int f(\mathbf{Y}|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})d\mathbf{Y}_1]_{\mathbf{Y}_2}=\infty\}$ be non-empty, and let $\mathscr{N}_2^c$ denote the set

complement of $\mathcal{N}_2$. Then, if $\mathbf{w} \sim \text{CM}(\mathbf{0}_{n,1}, \mathbf{I}_n, \boldsymbol{\alpha}, \boldsymbol{\kappa})$, a change of variables within the integral (see Proposition 1) gives,

$$\begin{aligned} 1 &= \int f(\mathbf{w}|\boldsymbol{\mu} = \mathbf{0}_{n,1}, \mathbf{V} = \mathbf{I}_n, \boldsymbol{\alpha}, \boldsymbol{\kappa})\, d\mathbf{w} = \int f(\mathbf{Y}|\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa})\, d\mathbf{Y} = \int\int f(\mathbf{Y}|\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) d\mathbf{Y}_1 d\mathbf{Y}_2 \\ &= \int_{\mathcal{N}_2} \int f(\mathbf{Y}|\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) d\mathbf{Y}_1 d\mathbf{Y}_2 + \int_{\mathcal{N}_2^c} \int f(\mathbf{Y}|\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) d\mathbf{Y}_1 d\mathbf{Y}_2 \\ &= \infty, \end{aligned}$$

which is a contradiction. Thus, we have that the conditional distribution of $\mathbf{Y}_1|\mathbf{Y}_2, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}$ is proper.

Proof of Theorem 2: Let $\mathbf{Y} \sim \text{CM}(\mathbf{V}\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}; \psi)$, and partition this n-dimensional random vector so that $\mathbf{Y} = (\mathbf{Y}_1', \mathbf{Y}_2')'$, where $\mathbf{Y}_1$ is r-dimensional and $\mathbf{Y}_2$ is $(n-r)$-dimensional. Additionally, denote the QR decomposition of $\mathbf{H}$ with

$$\mathbf{H} = \begin{bmatrix} \mathbf{Q}_1 & \mathbf{Q}_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 \\ \mathbf{0}_{n-r,n}, \end{bmatrix} \tag{A.2}$$

where recall that the $n \times r$ matrix $\mathbf{Q}_1$ satisfies $\mathbf{Q}_1'\mathbf{Q}_1 = \mathbf{I}_r$, the $n \times (n-r)$ matrix $\mathbf{Q}_2$ satisfies $\mathbf{Q}_2'\mathbf{Q}_2 = \mathbf{I}_{n-r}$ and $\mathbf{Q}_2'\mathbf{Q}_1 = \mathbf{0}_{n-r,r}$; $\mathbf{R}_1$ is a $r \times r$ upper triangular matrix; and $\sigma_2 > 0$. Then define $\mathbf{V}$ such that

$$\mathbf{V}^{-1} = \begin{bmatrix} \mathbf{Q}_1 & \mathbf{Q}_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 & \mathbf{0}_{r,n-r} \\ \mathbf{0}_{n-r,r} & \frac{1}{\sigma_2}\mathbf{I}_{n-r}, \end{bmatrix}. \tag{A.3}$$

Notice that

$$\mathbf{V} = \begin{bmatrix} (\mathbf{H}'\mathbf{H})^{-1}\mathbf{H}' \\ \sigma_2 \mathbf{Q}_2' \end{bmatrix}.$$

Thus, from (7) of the main text, $\mathbf{Y}$ can be written as

$$\begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{bmatrix} = \begin{bmatrix} (\mathbf{H}'\mathbf{H})^{-1}\mathbf{H}'\mathbf{w} \\ \sigma_2 \mathbf{Q}_2'\mathbf{w} \end{bmatrix}, \tag{A.4}$$

where the n-dimensional random vector $\mathbf{w} \sim \mathrm{CM}(\boldsymbol{\mu}, \mathbf{I}_n, \boldsymbol{\alpha}, \boldsymbol{\kappa};\ \psi)$. Multiplying both sides of (A.4) by $[\mathbf{I}_r, \mathbf{0}_{r,n-r}]$ we have

$$\mathbf{Y}_1 = (\mathbf{H}'\mathbf{H})^{-1}\mathbf{H}'\mathbf{w}, \tag{A.5}$$

and hence $(\mathbf{H}'\mathbf{H})^{-1}\mathbf{H}'\mathbf{w}$ is the marginal distribution associated with $\mathrm{CM}(\mathbf{V}\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa};\ \psi)$. To complete the proof, we need to find the distribution of $\mathbf{Y}_1$ and show that it is proportional to (11) of the main text. Our strategy will be to work with the conditional distribution from Theorem 1, and show that as σ_2 goes to infinity, the conditional distribution of $\mathbf{Y}_1$ given $\mathbf{Y}_2$ does not depend on $\mathbf{Y}_2$. This implies that the limit as σ_2 goes to infinity of the conditional distribution of $\mathbf{Y}_1$ given $\mathbf{Y}_2$ yields the marginal distribution of $\mathbf{Y}_1$ (see the argument in (A.9)).

For any $\mathbf{d} \in \mathbb{R}^n$, from Theorem 1, the conditional distribution of $\mathbf{Y}_1|\mathbf{Y}_2 = \mathbf{d}, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}$ has the following pdf:

$$f(\mathbf{Y}_1|\mathbf{Y}_2 = \mathbf{d}, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = \ M \exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi\left(\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\mu}^*\right)\right\},$$

where, $\boldsymbol{\mu}^* = \mathbf{V}^{-1}\mathbf{V}\boldsymbol{\mu} - \frac{1}{\sigma_2}\mathbf{Q}\mathbf{d} = \boldsymbol{\mu} - \frac{1}{\sigma_2}\mathbf{Q}\mathbf{d}$, and

$$M = \frac{\det(\mathbf{V}^{-1})\left\{\prod_{i=1}^{n} K(\alpha_i,\kappa_i)\right\}\exp\left(\frac{1}{\sigma_2}\boldsymbol{\alpha}'\mathbf{Q}_2\mathbf{d} - \boldsymbol{\alpha}'\boldsymbol{\mu}\right)}{\left[\int f\left(\mathbf{Y}|\mathbf{V}\boldsymbol{\mu}, \mathbf{V} = \left(\mathbf{H}\ \ \frac{1}{\sigma_2}\mathbf{Q}_2\right), \boldsymbol{\alpha}, \boldsymbol{\kappa}\right) d\mathbf{Y}_1\right]_{\mathbf{Y}_2=\mathbf{d}}}$$
$$= \frac{\exp\left(\frac{1}{\sigma_2}\boldsymbol{\alpha}'\mathbf{Q}_2\mathbf{d} - \boldsymbol{\alpha}'\boldsymbol{\mu}\right)}{\left[\int \exp\left\{\boldsymbol{\alpha}'\mathbf{V}^{-1}\mathbf{Y} - \boldsymbol{\alpha}'\boldsymbol{\mu} - \boldsymbol{\kappa}'\psi(\mathbf{V}^{-1}\mathbf{Y} - \boldsymbol{\mu})\right\} d\mathbf{Y}_1\right]_{\mathbf{Y}_2=\mathbf{d}}}.$$

This gives us that,

$$\lim_{\sigma_2\to\infty} f(\mathbf{Y}_1|\mathbf{Y}_2 = \mathbf{d}, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = \ M_1 \exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi\left(\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\mu}\right)\right\}, \tag{A.6}$$

where

$$M_1 = \frac{1}{\int \exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi(\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\mu})\right\} d\mathbf{Y}_1}, \tag{A.7}$$

and hence, the proportionality constant does not depend on $\mathbf{Y}_2$ (or $\mathbf{d}$) as σ_2 goes to infinity. Although we do not find the expression of M_1 we know that it is finite and non-zero. To see this, let $\mathbf{V}_c^{-1} = [\mathbf{H}, \mathbf{Q}_2]$, $\boldsymbol{\mu}_c \equiv \mathbf{V}_c(\boldsymbol{\mu} + \mathbf{Q}_2\mathbf{d}_f)$, and $g((\mathbf{Y}_1', \mathbf{Y}_2')'|\boldsymbol{\mu}_c, \mathbf{V}_c, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = \mathrm{CM}(\boldsymbol{\mu}_c, \mathbf{V}_c, \boldsymbol{\alpha}, \boldsymbol{\kappa})$. Then, for $(\mathbf{Y}_1', \mathbf{d}_c')' \in \mathscr{M}^n$,

$$\exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi\left(\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\mu}\right)\right\} \propto g\left(\mathbf{Y}_1|\mathbf{Y}_2 = \mathbf{d}_c, \boldsymbol{\mu}_c, \mathbf{V}_c, \boldsymbol{\alpha}, \boldsymbol{\kappa}\right). \tag{A.8}$$

From Theorem 1, we have that the conditional pdf $g\left(\mathbf{Y}_1|\mathbf{Y}_2 = \mathbf{d}_c, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa}\right)$ is proper. Consequently, Theorem 1, (A.7), and (A.8) imply that $\lim_{\sigma_2\to\infty} f(\mathbf{Y}_1|\mathbf{Y}_2 = \mathbf{d}, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa})$ integrates to 1.

Now, let

$$\rho(\mathbf{Y}_1, \boldsymbol{\mu}, \mathbf{H}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) \equiv M_1 \exp\left\{\boldsymbol{\alpha}'\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\kappa}'\psi\left(\mathbf{H}\mathbf{Y}_1 - \boldsymbol{\mu}\right)\right\},$$

so that $\rho(\mathbf{Y}_1, \mathbf{H}, \boldsymbol{\alpha}, \boldsymbol{\kappa}) = \lim_{\sigma_2\to\infty} f(\mathbf{Y}_1|\mathbf{Y}_2, \boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa})$. It follows from Equation (A.5) that $f(\mathbf{Y}_1|\boldsymbol{\mu}, \mathbf{V}, \boldsymbol{\alpha}, \boldsymbol{\kappa})$

does not depend on σ_2. This implies that

$$\begin{aligned}
f(\mathbf{Y}_1|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) &= \lim_{\sigma_2\to\infty} f(\mathbf{Y}_1|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) \\
&= \lim_{\sigma_2\to\infty} \int f(\mathbf{Y}_1|\mathbf{Y}_2,\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa})\, f(\mathbf{Y}_2|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) d\mathbf{Y}_2 \\
&= \int \lim_{\sigma_2\to\infty} f(\mathbf{Y}_1|\mathbf{Y}_2,\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) \lim_{\sigma_2\to\infty} f(\mathbf{Y}_2|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) d\mathbf{Y}_2 \\
&= \int \rho(\mathbf{Y}_1,\mathbf{H},\boldsymbol{\alpha},\boldsymbol{\kappa}) \lim_{\sigma_2\to\infty} f(\mathbf{Y}_2|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) d\mathbf{Y}_2 \\
&= \rho(\mathbf{Y}_1,\mathbf{H},\boldsymbol{\alpha},\boldsymbol{\kappa}) \lim_{\sigma_2\to\infty} \int f(\mathbf{Y}_2|\boldsymbol{\mu},\mathbf{V},\boldsymbol{\alpha},\boldsymbol{\kappa}) d\mathbf{Y}_2 \\
&= \rho(\mathbf{Y}_1,\boldsymbol{\mu},\mathbf{H},\boldsymbol{\alpha},\boldsymbol{\kappa}). \qquad \text{(A.9)}
\end{aligned}$$

Thus, $\rho(\mathbf{Y}_1,\mathbf{H},\boldsymbol{\alpha},\boldsymbol{\kappa})$ is proportional to the marginal pdf of $\mathbf{Y}_1$, which is the desired result.

Proof of Theorem 3: For $i=1,\ldots,n$, let $Z_i|\boldsymbol{\beta},\boldsymbol{\eta},\xi_i \overset{\text{ind}}{\sim} \text{EF}\left(\mathbf{x}_i'\boldsymbol{\beta}+\boldsymbol{\phi}_i'\boldsymbol{\eta}+\xi_i;\ \psi\right)$, and let $\boldsymbol{\beta} \sim \text{CM}\left(\mathbf{0}_{p,1},\mathbf{I}_p,\boldsymbol{\alpha}_\beta,\boldsymbol{\kappa};\ \psi\right)$, $\boldsymbol{\eta} \sim \text{CM}\left(\mathbf{0}_{r,1},\mathbf{V},\boldsymbol{\alpha}_\eta,\boldsymbol{\kappa}_\eta;\ \psi\right)$, and $\boldsymbol{\xi} \sim \text{MLG}\left(\mathbf{0}_{n,1},\mathbf{I}_n,\boldsymbol{\alpha}_\xi,\boldsymbol{\kappa}_\xi\right)$. Let $Y_i = \mathbf{x}_i'\boldsymbol{\beta}+\boldsymbol{\phi}_i\boldsymbol{\eta}+\xi_i$. The data model can be written as

$$\begin{aligned}
f(\mathbf{Z}|\cdot) &\underset{\mathbf{Z}}{\propto} \exp\left\{\sum_{i=1}^n Z_iY_i - \sum_{i=1}^n \psi(Y_i)\right\} = \exp\left\{\mathbf{Z}'\mathbf{Y} - \mathbf{J}_{n,1}'\psi(\mathbf{Y})\right\}, \qquad \text{(A.10)} \\
&= \exp\left\{\mathbf{Z}'\mathbf{X}\boldsymbol{\beta}+\mathbf{Z}'\boldsymbol{\Phi}\boldsymbol{\eta}+\mathbf{Z}'\boldsymbol{\xi} - \mathbf{J}_{n,1}'\psi(\mathbf{X}\boldsymbol{\beta}+\boldsymbol{\Phi}\boldsymbol{\eta}+\boldsymbol{\xi})\right\}, \qquad \text{(A.11)}
\end{aligned}$$

where $\mathbf{Y} \equiv (Y_1,\ldots,Y_n)'$, the $n\times p$ matrix $\mathbf{X} \equiv (\mathbf{x}_1,\ldots,\mathbf{x}_n)'$, the $n\times r$ matrix $\Phi \equiv (\boldsymbol{\phi}_1,\ldots,\boldsymbol{\phi}_n)'$, and $\underset{\mathbf{Z}}{\propto}$ denotes the "proportional to as a function of $\mathbf{Z}$" symbol. It follows that

$$f(\mathbf{Z}|\cdot) \underset{\boldsymbol{\beta}}{\propto} \exp\left(\mathbf{Z}'\mathbf{X}\boldsymbol{\beta} - b\mathbf{J}_{n,1}'\psi\left(\mathbf{X}\boldsymbol{\beta}+\boldsymbol{\Phi}\boldsymbol{\eta}+\boldsymbol{\xi}\right)\right) h \qquad \text{(A.12)}$$

$$f(\mathbf{Z}|\cdot) \underset{\boldsymbol{\eta}}{\propto} \exp\left(\mathbf{Z}'\boldsymbol{\Phi}\boldsymbol{\eta} - b\mathbf{J}_{n,1}'\psi\left(\boldsymbol{\Phi}\boldsymbol{\eta}+\mathbf{X}\boldsymbol{\beta}+\boldsymbol{\xi}\right)\right) h \qquad \text{(A.13)}$$

$$f(\mathbf{Z}|\cdot) \underset{\boldsymbol{\xi}}{\propto} \exp\left(\mathbf{Z}'\boldsymbol{\xi} - b\mathbf{J}_{n,1}'\psi\left(\boldsymbol{\xi}+\mathbf{X}\boldsymbol{\beta}+\boldsymbol{\Phi}\boldsymbol{\eta}\right)\right) h, \qquad \text{(A.14)}$$

where $h = \{\prod_{i=1}^{n} I(\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\}$. From Proposition 1($i$) we have that

$$f(\boldsymbol{\beta}|\mathbf{V}_\beta, \boldsymbol{\alpha}_\beta, \boldsymbol{\kappa}_\beta) \propto \exp\left\{\boldsymbol{\alpha}_\beta'\mathbf{V}_\beta^{-1}\boldsymbol{\beta} - \boldsymbol{\kappa}_\beta'\psi\left(\mathbf{V}_\beta^{-1}\boldsymbol{\beta}\right)\right\}, \tag{A.15}$$

$$f(\boldsymbol{\eta}|\mathbf{V}, \boldsymbol{\alpha}_\eta, \boldsymbol{\kappa}_\eta) \propto \exp\left\{\boldsymbol{\alpha}_\eta'\mathbf{V}^{-1}\boldsymbol{\eta} - \boldsymbol{\kappa}_\eta'\psi\left(\mathbf{V}^{-1}\boldsymbol{\eta}\right)\right\}, \tag{A.16}$$

$$f(\boldsymbol{\xi}|\mathbf{V}_\xi \boldsymbol{\alpha}_\xi, \boldsymbol{\kappa}_\xi) \propto \exp\left\{\boldsymbol{\alpha}_\xi'\mathbf{V}_\xi^{-1}\boldsymbol{\xi} - \boldsymbol{\kappa}_\xi'\psi\left(\mathbf{V}_\xi^{-1}\boldsymbol{\xi}\right)\right\}. \tag{A.17}$$

Then, define the $(n+p)\times p$, $(n+r)\times r$, and $2n\times n$ matrices

$$\mathbf{H}_\beta \equiv \begin{bmatrix} \mathbf{X} \\ \mathbf{V}_\beta^{-1} \end{bmatrix}, \qquad \mathbf{H}_\eta \equiv \begin{bmatrix} \boldsymbol{\Phi} \\ \mathbf{V}^{-1} \end{bmatrix}, \qquad \text{and } \mathbf{H}_\xi \equiv \begin{bmatrix} \mathbf{I}_n \\ \mathbf{V}_\xi^{-1} \end{bmatrix}.$$

Using (A.12) and (A.15) we have that

$$\begin{aligned} f(\boldsymbol{\beta}|\cdot) &\underset{\boldsymbol{\beta}}{\propto} f(\mathbf{Z}|\cdot) f(\boldsymbol{\beta}|\boldsymbol{\alpha}_\beta, \boldsymbol{\kappa}_\beta) \underset{\boldsymbol{\beta}}{\propto} \exp\left\{\mathbf{Z}'\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\alpha}_\beta'\mathbf{V}_\beta^{-1}\boldsymbol{\beta} - \boldsymbol{\kappa}_\beta'\psi\left(\mathbf{V}_\beta^{-1}\boldsymbol{\beta}\right) - b\mathbf{J}_{n,1}'\psi(\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\Phi}\boldsymbol{\eta} + \boldsymbol{\xi})\right\} h \\ &= \exp\left\{\mathbf{Z}'\mathbf{X}\boldsymbol{\beta} + \mathbf{d}_\beta'\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\alpha}_\beta'\mathbf{V}_\beta^{-1}\boldsymbol{\beta} - \mathbf{d}_\beta'\mathbf{X}\mathbf{V}_\beta\mathbf{V}_\beta^{-1}\boldsymbol{\beta} - \boldsymbol{\kappa}_\beta'\psi\left(\mathbf{V}_\beta^{-1}\boldsymbol{\beta}\right) - b\mathbf{J}_{n,1}'\psi(\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\Phi}\boldsymbol{\eta} + \boldsymbol{\xi})\right\} h \\ &= \exp\left\{(\mathbf{Z}' + \mathbf{d}_\beta', \boldsymbol{\alpha}_\beta' - \mathbf{d}_\beta'\mathbf{X}\mathbf{V}_\beta)\mathbf{H}_\beta\boldsymbol{\beta} - (b\mathbf{J}_{n,1}', \boldsymbol{\kappa}_\beta')\psi\left(\mathbf{H}_\beta\boldsymbol{\beta} - \boldsymbol{\mu}_\beta^*\right)\right\} h \\ &\propto \mathrm{CM_c}\left\{\boldsymbol{\mu}_\beta^*, \mathbf{H}_\beta, (\mathbf{Z}' + \mathbf{d}_\beta', \boldsymbol{\alpha}_\beta' - \mathbf{d}_\beta'\mathbf{X}\mathbf{V}_\beta)', (b\mathbf{J}_{n,1}', \boldsymbol{\kappa}_\beta')';\ \psi\right\}\left\{\prod_{i=1}^{n} I(\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\right\}, \end{aligned} \tag{A.18}$$

where $\boldsymbol{\mu}_\beta^* \equiv \left(-\boldsymbol{\eta}'\boldsymbol{\Phi}' - \boldsymbol{\xi}', \mathbf{0}_{1,p}\right)'$, and

$$\mathbf{d}_\beta \equiv \frac{\min(\boldsymbol{\alpha}_\beta)}{\min(\boldsymbol{\alpha}_\beta) + \max\left\{\mathrm{abs}\left(\mathbf{J}_{n,1}'\mathbf{X}\mathbf{V}_\beta\right)\right\}}\mathbf{e}_0(\mathbf{Z}) - \frac{\min(\boldsymbol{\kappa}_\beta - \boldsymbol{\alpha}_\beta)}{\min(\boldsymbol{\kappa}_\beta - \boldsymbol{\alpha}_\beta) + \max\left\{\mathrm{abs}\left(\mathbf{J}_{n,1}'\mathbf{X}\mathbf{V}_\beta\right)\right\}}\mathbf{e}_b(\mathbf{Z}), \tag{A.19}$$

and max(**k**) and min(**k**) gives the maximum and minimum values of the elements in the m-dimensional real-valued vector **k**. Let $\mathbf{e}_0(\mathbf{Z}) \equiv (I(Z_1 = 0), \ldots, I(Z_n = 0))'$ if Z_i is bounded below by zero, and set $\mathbf{e}_0(\mathbf{Z}) \equiv \mathbf{0}_{n,1}$ otherwise. Let $\mathbf{e}_b(\mathbf{Z}) \equiv (I(Z_1 = b), \ldots, I(Z_n = b))'$ if Z_i is bounded

above by b, and set $\mathbf{e}_b(\mathbf{Z}) \equiv \mathbf{0}_{n,1}$ otherwise. The vector $\mathbf{d}_\beta$ accommodates values of $\mathbf{Z}$ on the boundary. For example, the log gamma distribution the shape and scale parameters need to be nonnegative, and for the logit beta distribution the shape parameter must be positive and less than the scale parameter. The vector $\mathbf{d}_\beta$ adjusts the data vector $\mathbf{Z}$, but this is not an artificial translation, which are sometimes applied to non-Gaussian settings (Bishop et al., 1977). This can be seen in the second line of (A.18), where the term $\mathbf{d}_\beta$ is introduced through $\mathbf{d}_\beta'\mathbf{X}\boldsymbol{\beta} - \mathbf{d}_\beta'\mathbf{X}\boldsymbol{\beta} = \mathbf{0}_{1,p}$, and hence does not change the value of the full-conditional distribution for a given set of values for the data, process, and parameters.

In a similar manner we can find the full conditional distributions associated with $\boldsymbol{\eta}$ and $\boldsymbol{\xi}$. Using (A.13) and (A.16) we have that

$$\begin{aligned}
f(\boldsymbol{\eta}|\cdot) &\underset{\boldsymbol{\eta}}{\propto} f(\mathbf{Z}|\cdot) f(\boldsymbol{\eta}|\boldsymbol{\alpha}_\beta, \boldsymbol{\kappa}_\beta) \underset{\boldsymbol{\eta}}{\propto} \exp\left\{\mathbf{Z}'\boldsymbol{\Phi}\boldsymbol{\eta} + \boldsymbol{\alpha}_\eta'\mathbf{V}^{-1}\boldsymbol{\eta} - \boldsymbol{\kappa}_\eta'\psi\left\{\mathbf{V}^{-1}\boldsymbol{\eta}\right\} - b\mathbf{J}_{n,1}'\psi(\boldsymbol{\Phi}\boldsymbol{\eta} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\xi})\right\} h \\
&= \exp\left\{\mathbf{Z}'\boldsymbol{\Phi}\boldsymbol{\eta} + \mathbf{d}_\eta'\boldsymbol{\Phi}\boldsymbol{\eta} + \boldsymbol{\alpha}_\eta'\mathbf{V}^{-1}\boldsymbol{\eta} - \mathbf{d}_\eta'\boldsymbol{\Phi}\mathbf{V}\mathbf{V}^{-1}\boldsymbol{\eta} - \boldsymbol{\kappa}_\eta'\psi\left\{\mathbf{V}^{-1}\boldsymbol{\eta}\right\} - b\mathbf{J}_{n,1}'\psi(\boldsymbol{\Phi}\boldsymbol{\eta} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\xi})\right\} h \\
&= \exp\left\{(\mathbf{Z}' + \mathbf{d}_\eta', \boldsymbol{\alpha}_\eta' - \mathbf{d}_\eta'\boldsymbol{\Phi}\mathbf{V})\mathbf{H}_\eta\boldsymbol{\eta} - (b\mathbf{J}_{r,1}', \boldsymbol{\kappa}_\eta')\psi\left(\mathbf{H}_\eta\boldsymbol{\eta} - \boldsymbol{\mu}_\eta^*\right)\right\} h \\
&\propto \mathrm{CM_c}\left\{\boldsymbol{\mu}_\eta^*, \mathbf{H}_\eta, (\mathbf{Z}' + \mathbf{d}_\eta', \boldsymbol{\alpha}_\eta' - \mathbf{d}_\eta'\boldsymbol{\Phi}\mathbf{V})', (b\mathbf{J}_{r,1}', \boldsymbol{\kappa}_\eta')';\ \psi\right\}\left\{\prod_{i=1}^{n} I(\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\right\},
\end{aligned} \tag{A.20}$$

where $\boldsymbol{\mu}_\eta^* \equiv \left(-\boldsymbol{\beta}'\mathbf{X}' - \boldsymbol{\xi}', \mathbf{0}_{1,r}\right)'$, and

$$\mathbf{d}_\eta \equiv \frac{\min(\boldsymbol{\alpha}_\eta)}{\min(\boldsymbol{\alpha}_\eta) + \max\left\{\mathrm{abs}\left(\mathbf{J}_{n,1}'\boldsymbol{\Phi}\mathbf{V}\right)\right\}}\mathbf{e}_0(\mathbf{Z}) - \frac{\min(\boldsymbol{\kappa}_\eta - \boldsymbol{\alpha}_\eta)}{\min(\boldsymbol{\kappa}_\eta - \boldsymbol{\alpha}_\eta) + \max\left\{\mathrm{abs}\left(\mathbf{J}_{n,1}'\boldsymbol{\Phi}\mathbf{V}\right)\right\}}\mathbf{e}_b(\mathbf{Z}).$$

Finally, using (A.14) and (A.17) we have that

$$
\begin{aligned}
f(\boldsymbol{\xi}|\cdot) &\underset{\boldsymbol{\xi}}{\propto} f(\mathbf{Z}|\cdot) f(\boldsymbol{\xi}|\boldsymbol{\alpha}_\beta, \boldsymbol{\kappa}_\beta) \underset{\boldsymbol{\beta}}{\propto} \exp\left\{\mathbf{Z}'\boldsymbol{\xi} + \boldsymbol{\alpha}_\xi' \mathbf{V}_\xi^{-1}\boldsymbol{\xi} - \boldsymbol{\kappa}_\xi' \psi\left(\mathbf{V}_\xi^{-1}\boldsymbol{\xi}\right) - b\mathbf{J}_{n,1}'\psi(\boldsymbol{\xi} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\Phi}\boldsymbol{\eta})\right\} h \\
&= \exp\left\{\mathbf{Z}'\boldsymbol{\xi} + \mathbf{d}_\xi'\boldsymbol{\xi} + \boldsymbol{\alpha}_\xi' \mathbf{V}_\xi^{-1}\boldsymbol{\xi} - \mathbf{d}_\xi'\mathbf{V}_\xi\mathbf{V}_\xi^{-1}\boldsymbol{\xi} - \boldsymbol{\kappa}_\xi' \psi\left(\mathbf{V}_\xi^{-1}\boldsymbol{\xi}\right) - b\mathbf{J}_{n,1}'\psi(\boldsymbol{\xi} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\Phi}\boldsymbol{\eta})\right\} h \\
&= \exp\left\{(\mathbf{Z}' + \mathbf{d}_\xi', \boldsymbol{\alpha}_\xi' - \mathbf{d}_\xi'\mathbf{V}_\xi)\mathbf{H}_\xi\boldsymbol{\xi} - (b\mathbf{J}_{n,1}', \boldsymbol{\kappa}_\xi')\psi\left(\mathbf{H}_\xi\boldsymbol{\xi} - \boldsymbol{\mu}_\xi^*\right)\right\} h \\
&\propto \mathrm{CM_c}\left\{\boldsymbol{\mu}_\xi^*, \mathbf{H}_\xi, (\mathbf{Z}' + \mathbf{d}_\xi', \boldsymbol{\alpha}_\xi' - \mathbf{d}_\xi')', (b\mathbf{J}_{n,1}', \boldsymbol{\kappa}_\xi')';\ \psi\right\}\left\{\prod_{i=1}^n I(\mathbf{x}_i'\boldsymbol{\beta} + \boldsymbol{\psi}_i'\boldsymbol{\eta} + \xi_i \in \mathscr{Y})\right\},
\end{aligned}
\tag{A.21}
$$

where $\boldsymbol{\mu}_\xi^* \equiv \left(-\boldsymbol{\beta}'\mathbf{X}' - \boldsymbol{\eta}'\boldsymbol{\Phi}', \mathbf{0}_{1,n}\right)'$, and

$$
\mathbf{d}_\xi \equiv \frac{\min(\boldsymbol{\alpha}_\xi)}{\min(\boldsymbol{\alpha}_\xi) + \max(\mathbf{J}_{n,1}'\mathbf{V}_\xi)}\mathbf{e}_0(\mathbf{Z}) - \frac{\min(\boldsymbol{\kappa}_\xi - \boldsymbol{\alpha}_\xi)}{\min(\boldsymbol{\kappa}_\xi - \boldsymbol{\alpha}_\xi) + \max(\mathbf{J}_{n,1}'\mathbf{V}_\xi)}\mathbf{e}_b(\mathbf{Z}).
$$

The important quantities from this proof are organized into Table 2.

Proof of Theorem 4: Using induction we find that,

$$
\begin{aligned}
f(\mathbf{Y}|\boldsymbol{\mu} = \mathbf{0}_{n,1}, \mathbf{V}, \boldsymbol{\alpha}_Y, \boldsymbol{\kappa}_Y) &\underset{\mathbf{V}}{\propto} \exp\left[\boldsymbol{\alpha}_Y'\mathbf{V}^{-1}\mathbf{Y} - \boldsymbol{\kappa}_Y'\psi\left\{\mathbf{V}^{-1}\mathbf{Y}\right\}\right] \\
&\propto \exp\left[\sum_{i=2}^n \sum_{j=1}^{i-1} \alpha_{Y,i} v_{i,j} Y_j - \sum_{i=2}^n \kappa_{Y,i}\psi\left(\sum_{j=1}^{i-1} Y_j v_{i,j} + Y_i\right)\right] \\
&= \prod_{i=2}^n \exp\left\{\alpha_i \boldsymbol{\Sigma}_i \mathbf{v}_i - \kappa_i \psi(\boldsymbol{\Sigma}_i \mathbf{v}_i + Y_i)\right\}
\end{aligned}
$$

where $\boldsymbol{\Sigma}_i' = \left(Y_j : j = 1, \ldots, i-1\right)'$. Thus, the full conditional distribution is given by

$$
\begin{aligned}
f(\mathbf{v}_2, \ldots, \mathbf{v}_r|\cdot) &\underset{\mathbf{V}}{\propto} f(\mathbf{Y}|\boldsymbol{\mu} = \mathbf{0}_{n,1}, \mathbf{V}, \boldsymbol{\alpha}_Y, \boldsymbol{\kappa}_Y) \prod_{i=2}^{n} f(\mathbf{v}_i) \\
&\underset{\mathbf{V}}{\propto} \prod_{i=2}^{n} \exp\left[(\alpha_{Y,i}, \boldsymbol{\alpha}_i')\mathbf{H}_{\gamma,i}(\boldsymbol{\gamma}_i - \boldsymbol{\mu}_{\gamma,i}^*) - (\kappa_{Y,i}, \boldsymbol{\kappa}_i')\psi\left\{\mathbf{H}_{\gamma,i}\boldsymbol{\gamma}_i - \boldsymbol{\mu}_{\gamma,i}\right\}\right], \\
&\underset{\mathbf{V}}{\propto} \prod_{i=2}^{n} \mathrm{CM_c}\left(\boldsymbol{\mu}_{\gamma,i}, \mathbf{H}_{\gamma,i}, (\alpha_{Y,i}, \boldsymbol{\alpha}_i)', (\kappa_{Y,i}, \boldsymbol{\kappa}_i')'; \psi\right),
\end{aligned}
$$

where

$$
\begin{aligned}
\mathbf{H}_{\gamma,i} &\equiv \begin{bmatrix} \boldsymbol{\Sigma}_i \\ \\ \mathbf{C}_i \end{bmatrix}, \\
\boldsymbol{\mu}_{\gamma,i} &= (Y_i, \ \mathbf{0}_{1,i-1})',
\end{aligned}
$$

which is the desired result.

Proof of Theorem 5: Consider the transformation $Q = \left(\frac{\psi''(0)}{\psi'(0)}\right)^{1/2} \alpha^{1/2} W$, where $W \sim \mathrm{DY}\left(\alpha, \frac{\alpha}{\psi'(0)}; \psi\right)$. Then we have that

$$
f(Q|\alpha, \kappa) \propto \exp\left[\left(\frac{\psi'(0)}{\psi''(0)}\right)^{1/2} \alpha^{1/2} Q - \frac{\alpha}{\psi'(0)}\, \psi\left\{\left(\frac{\psi'(0)}{\psi''(0)}\right)^{1/2} \alpha^{-1/2} Q\right\}\right],
$$

and using the Taylor Series expansion of $\exp(x)$ we have

$$
\begin{aligned}
&f(Q|\alpha,\kappa) \\
&\propto \exp\left[\left(\frac{\psi'(0)}{\psi''(0)}\right)^{1/2}\alpha^{1/2}Q\right. \\
&\left.-\frac{\alpha}{\psi'(0)}\left\{\psi'(0)\left(\frac{\psi'(0)}{\psi''(0)}\right)^{1/2}\alpha^{-1/2}Q+\psi''(0)\left(\frac{\psi'(0)}{\psi''(0)}\right)\alpha^{-1}\frac{Q^2}{2}+O\left(\frac{\psi''(0)^{3/2}}{\psi'(0)^{3/2}}\alpha^{-3/2}Q^3\right)\right\}\right],
\end{aligned}
$$

where "$O(\cdot)$" is the "Big-O" notation (e.g., see Lehmann, 1999, among others). Then, letting α go to infinity yields,

$$\lim_{\alpha\to\infty} f(Q|\alpha,\kappa) \propto \exp\left(-\frac{Q^2}{2}\right) \propto \text{Normal}(0,1).$$

Thus, Q converges in distribution to a standard normal distribution as α goes to infinity. Now suppose $\mathbf{w} = (w_1,\ldots,w_n)' \sim \text{CM}\left(\mathbf{0}_n, \alpha^{1/2}\mathbf{I}_n, \alpha\mathbf{J}_{n,1}, \frac{\alpha}{\psi'(0)}\mathbf{J}_{n,1};\ \psi\right)$. Then it follows from the result above that $\left(\frac{\alpha}{\psi'(0)}\right)^{1/2}\mathbf{w}$ converges to a standard multivariate Gaussian distribution. Now, define the transformation $\mathbf{Y} = \boldsymbol{\mu} + \mathbf{V}(\alpha^{1/2}\mathbf{w})$, which follows a $\text{CM}(\boldsymbol{\mu}, \alpha^{1/2}\mathbf{V}, \alpha\mathbf{J}_{n,1}, \frac{\alpha}{\psi'(0)}\mathbf{J}_{n,1};\ \psi)$. It follows from Theorem 5.1.8 of Lehmann (1999), and the fact that $\frac{\alpha}{\psi'(0)}\mathbf{w}$ converges to a standard Gaussian distribution, that $\mathbf{Y}$ converges in distribution to a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{V}\mathbf{V}'$.

6 Appendix B: The Gibbs Sampler

The Gibbs sampler associated with (20) of the main text requires one to compute certain quantities. These values are listed in Table 3. To aid the reader, we provide step-by-step instructions for implementing the Gibbs sampler associated with (20) of the main text as follows.

1. Initialize $\boldsymbol{\beta}$, $\boldsymbol{\eta}$, $\boldsymbol{\xi}$, $\{\mathbf{v}_i\}$, $\boldsymbol{\alpha}_\eta = (\alpha_{\eta,1},\ldots,\alpha_{\eta,r})'$, α_ξ, $\boldsymbol{\kappa}_\eta = (\kappa_{\eta,1},\ldots,\kappa_{\eta,r})'$, and κ_ξ. De-

note these initializations with $\boldsymbol{\beta}^{[0]}$, $\boldsymbol{\eta}^{[0]}$, $\boldsymbol{\xi}^{[0]}$, $\{\mathbf{v}_i^{[0]}\}$, $\boldsymbol{\alpha}_\eta^{[0]} = (\alpha_{\eta,1}^{[0]}, \ldots, \alpha_{\eta,r}^{[0]})'$, $\alpha_\xi^{[0]}$, $\boldsymbol{\kappa}_\eta^{[0]} = (\kappa_{\eta,1}^{[0]}, \ldots, \kappa_{\eta,r}^{[0]})'$, and $\kappa_\xi^{[0]}$. Set $m = 1$.

2. Set $\boldsymbol{\beta}^{[m]}$ equal to a draw from $\text{CM}_\text{c}(\boldsymbol{\mu}_\beta, \mathbf{H}_\beta, \boldsymbol{\alpha}_\beta^*, \boldsymbol{\kappa}_\beta^*)$ using Theorem 2, where $\boldsymbol{\mu}_\beta$, $\mathbf{H}_\beta$, $\boldsymbol{\alpha}_\beta^*$, and $\boldsymbol{\kappa}_\beta^*$ are computed using Table 2 and the values of $\boldsymbol{\eta}^{[b-1]}$ and $\boldsymbol{\xi}^{[b-1]}$.

3. Set $\boldsymbol{\eta}^{[m]}$ equal to a draw from $\text{CM}_\text{c}(\boldsymbol{\mu}_\eta, \mathbf{H}_\eta, \boldsymbol{\alpha}_\eta^*, \boldsymbol{\kappa}_\eta^*)$ using Theorem 2, where $\boldsymbol{\mu}_\eta$, $\mathbf{H}_\eta$, $\boldsymbol{\alpha}_\eta^*$, and $\boldsymbol{\kappa}_\eta^*$ are computed using Table 2 and the values of $\boldsymbol{\beta}^{[m]}$, $\boldsymbol{\xi}^{[b-1]}$, $\boldsymbol{\alpha}_\eta^{[b-1]}$, $\boldsymbol{\kappa}_\eta^{[b-1]}$, and $\{\mathbf{v}_i^{[b-1]}\}$.

4. Set $\boldsymbol{\xi}^{[m]}$ equal to a draw from $\text{CM}_\text{c}(\boldsymbol{\mu}_\xi, \mathbf{H}_\xi, \boldsymbol{\alpha}_\xi^*, \boldsymbol{\kappa}_\xi^*)$ using Theorem 2, $\mathbf{H}_\xi$, $\boldsymbol{\alpha}_\xi^*$, and $\boldsymbol{\kappa}_\xi^*$ are computed using Table 2 and the values of $\boldsymbol{\eta}^{[m]}$, $\boldsymbol{\beta}^{[m]}$, $\alpha_\xi^{[b-1]}$, and $\kappa_\xi^{[b-1]}$.

5. For $i = 2, \ldots, r$ set $\mathbf{v}_i^{[m]}$ equal to a value generated from $\text{CM}_\text{c}(\boldsymbol{\mu}_{\gamma,i}, \mathbf{H}_{\gamma,i}, \boldsymbol{\alpha}_{\gamma,i}, \boldsymbol{\kappa}_{\gamma,i})$, where $\boldsymbol{\mu}_{\gamma,i}, \mathbf{H}_{\gamma,i}, \boldsymbol{\alpha}_{\gamma,i}$, and $\boldsymbol{\kappa}_{\gamma,i}$ are given in Table 3, and are computed using $\boldsymbol{\eta}^{[m]}$, $\alpha_{\eta,i}^{[b-1]}$, and $\kappa_{\eta,i}^{[b-1]}$.

6. For each $i = 1, \ldots, r$, use a slice sampler to set $\alpha_{\eta,i}^{[m]}$ and $\kappa_{\eta,i}^{[m]}$ to a value generated from the pdf:

$$f(\alpha_{\eta,i}, \kappa_{\eta,i} | \cdot) \propto \exp\left[(\gamma_{\eta,1} + \mathbf{e}_i' \mathbf{V}^{-1} \boldsymbol{\eta}) \alpha_{\eta,i} + \left\{ \gamma_{\eta,2} - \mathbf{e}_i' \boldsymbol{\psi}(\mathbf{V}^{-1} \boldsymbol{\eta}) \right\} \kappa_{\eta,i} - (\rho_\eta + 1) \log\left\{ \frac{1}{K(\alpha_{\eta,i}, \kappa_{\eta,i})} \right\} \right].$$

7. Using a slice sampler to set $\alpha_\xi^{[m]}$ and $\kappa_\xi^{[m]}$ equal to values generated from the pdf:

$$f(\alpha_\xi, \kappa_\xi | \cdot) \propto \exp\left[(\gamma_{\xi,1} + \mathbf{J}_{1,n} \boldsymbol{\xi}) \alpha_\xi + \left\{ \gamma_{\xi,2} - \mathbf{J}_{1,n} \psi(\boldsymbol{\xi}) \right\} \kappa_\xi - (\rho_\xi + 1) \log\left\{ \frac{1}{K(\alpha_\xi, \kappa_\xi)} \right\} \right].$$

8. Set $m = m + 1$

9. Repeat steps 2 through 8 until convergence of the Gibbs sampler.

It is straightforward to adjust this Gibbs sampler in variety of ways to be more appropriate for a particular problem. For example, one could consider different hyperparameters, different basis functions $\{\boldsymbol{\phi}_j\}$, assume constant DY parameters associated with $\boldsymbol{\eta}$, and assume heterogeneous DY parameters associated with $\boldsymbol{\beta}$ and $\boldsymbol{\xi}$.

It is important to note that many software packages have built in functions to simulate from beta and gamma distributions, which are needed when $j = k = 2$ and $j = k = 3$, respectively. However, it is common for the Gibbs sampler to produce small values of shape and scale parameters, which may lead to computational errors when simulating from a beta or a gamma distribution. In this setting, we simulate beta and gamma random variables using strategies outlined in Devroye (1986, pgs. 181, 182, and 419). Additionally, if the shape and scale parameters are so small (i.e., close to zero) that it is not possible to simulate the beta and gamma random variables using the techniques in Devroye (1986), we reject the proposed sample. However, after a sufficient burn-in period of the Gibbs sampler, the acceptance rate is approximately equal to one.

References

Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2015). *Hierarchical Modeling and Analysis for Spatial Data*. London, UK: Chapman and Hall.

Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008). "Gaussian predictive process models for large spatial data sets." *Journal of the Royal Statistical Society, Series B*, 70, 825–848.

Bishop, Y. M., Fienberg, S. E., and Holland, P. W. (1977). *Discrete Multivariate Analysis: Theory and Practice*. New York, NY: Springer-Verlag.

Bradley, J., Holan, S., and Wikle, C. (2015a). "Computationally Efficient Distribution Theory for Bayesian Inference of High-Dimensional Dependent Count-Valued Data." *arXiv preprint: 1512.07273*.

— (2015b). "Multivariate spatio-temporal models for high-dimensional areal data with application to Longitudinal Employer-Household Dynamics." *The Annals of Applied Statistics*, 9, 1761–1791.

Bradley, J., Wikle, C., and Holan, S. (2015c). "Bayesian spatial change of support for count-valued survey data." *Journal of the American Statistical Association*, forthcoming.

Casella, G. and Berger, R. (2002). *Statistical Inference*. Pacific Grove, CA: Duxbury.

Castruccio, S., Ombao, H., and Genton, M. G. (2016). "A Scalable Multi-Resolution Spatio-Temporal Model for Brain Activation and Connectivity in fMRI Data." *arXiv preprint: 1602.02435*.

Chen, Z. and Dunson, D. B. (2003). "Random effects selection in linear mixed models." *Biometrics*, 59, 762–769.

Conway, R. W. and Maxwell, W. L. (1962). "A queuing model with state dependent service rates." *Journal of the American Statistical Association*, 12, 132–136.

Cox, T. F. (2005). *An Introduction to Multivariate Data Analysis*. London: Hodder Arnold.

Cressie, N. and Johannesson, G. (2006). "Spatial prediction for massive data sets." In *Australian Academy of Science Elizabeth and Frederick White Conference*, 1–11. Australian Academy of Science, Canberra.

— (2008). "Fixed rank kriging for very large spatial data sets." *Journal of the Royal Statistical Society, Series B*, 70, 209–226.

Cressie, N., Shi, T., and Kang, E. L. (2010a). " Fixed Rank Filtering for spatio-temporal data." *Journal of Computational and Graphical Statistics*, 19, 724–745.

— (2010b). " Using temporal variability to improve spatial mapping with application to satellite data." *Canadian Journal of Statistics*, 38, 271–289.

Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Hoboken, NJ: Wiley.

Daniels, M. J. and Pourahmadi, M. (2002). "Dynamic models and Bayesian analysis of covariance matrices in longitudinal data." *Biometrika*, 89, 553–566.

De Oliveira, V. (2013). "Hierarchical Poisson models for spatial count data." *Journal of Multivariate Analysis*, 122, 393–408.

Demirhan, H. and Hamurkaroglu, C. (2011). "On a multivariate log-gamma distribution and the use of the distribution in the Bayesian analysis." *Journal of Statistical Planning and Inference*, 141, 1141–1152.

Devroye, L. (1986). *Non-Uniform Random Variate Generation*. New York, NY: Springer-Verlag.

Diaconis, P. and Ylvisaker, D. (1979). "Conjugate priors for exponential families." *The Annals of Statistics*, 17, 269–281.

Diggle, P. J., Tawn, J. A., and Moyeed, R. A. (1998). "Model-based geostatistics." *Journal of the Royal Statistical Society, Series C*, 47, 299–350.

Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). "Hybrid Monte Carlo." *Physics Letters B*, 195, 216–222.

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). "Least angle regression." *Annals of Statistics*, 32, 407–499.

Everitt, B. and Hothorn, T. (2011). *An Introduction to Applied Multivariate Analysis with R*. New York: Springer.

Finley, A. O., Sang, H., Banerjee, S., and Gelfand, A. E. (2009). "Improving the performance of predictive process modeling for large datasets." *Computational Statistics and Data Analysis*, 53, 2873–2884.

Gelman, A. (2006). "Prior distributions for variance parameters in hierarchical models." *Bayesian Analysis*, 1, 515–533.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis, 3rd edn.*. Boca Raton, FL: Chapman and Hall/CRC.

Griffith, D. (2000). "A linear regression solution to the spatial autocorrelation problem." *Journal of Geographical Systems*, 2, 141–156.

— (2002). "A spatial filtering specification for the auto-Poisson model." *Statistics and Probability Letters*, 58, 245–251.

— (2004). "A spatial filtering specification for the auto-logistic model." *Environment and Planning A*, 36, 1791–1811.

Holan, S. H. and Wikle, C. K. (2012). "Semiparametric Dynamic Design of Monitoring Networks for Non-Gaussian Spatio-Temporal Data." In *Spatio-Temporal Design Advances in Efficient Data Acquisition*, eds. J. Mateu and W. Muller. New York, NY: Wiley.

— (2016). "Hierarchical dynamic generalized linear mixed models for discrete-valued spatio-temporal data." In *Handbook of Discrete–Valued Time Series*. R. A. Davis, S. H. Holan, R. Lund, and N. Ravishanker (eds). CRC Press.

Hooten, M. B., Larsen, D. R., and Wikle, C. K. (2003). "Predicting the spatial distribution of ground flora on large domains using a hierarchical Bayesian model." *Landscape Ecology*, 18, 487–502.

Hughes, J. and Haran, M. (2013). "Dimension reduction and alleviation of confounding for spatial generalized linear mixed model." *Journal of the Royal Statistical Society, Series B*, 75, 139–159.

Jolliffe, I. T. (2002). *Principal Components Analysis, 2nd edn.*. New York: Springer Verlag.

Kang, E. L. and Cressie, N. (2011). "Bayesian inference for the spatial random effects model." *Journal of the American Statistical Association*, 106, 972 – 983.

Katzfuss, M. and Cressie, N. (2011). "Spatio-temporal smoothing and EM estimation for massive remote-sensing data sets." *Journal of Time Series Analysis*, 32, 430–446.

— (2012). "Bayesian hierarchical spatio-temporal smoothing for very large datasets." *Environmetrics*, 23, 94–107.

Kotz, S., Balakrishnan, N., and Johnson, N. (2000). *Continuous Multivariate Distributions, Volume 1: Models and Applications*. New York, NY: Wiley.

Lambert, D. (2006). "Zero-Inflated Poisson Regression, with an Application to Defects in Manufacturing." *Technometrics*, 34, 1–14.

Lange, K., Papp, J. C., Sinsheimer, J. S., and Sobel, E. M. (2014). "Next-Generation Statistical Genetics: Modeling, Penalization, and Optimization in High-Dimensional Data." *Annual Review of Statistics and Its Application*, 1, 279–300.

Lee, Y. and Nelder, J. A. (1974). "Double hierarchical generalized linear models with discussion." *Applied Statistics*, 55, 129–185.

— (1996). "Hierarchical generalized linear models (with discussion)." *Journal of the Royal Statistical Society, Series B*, 58, 619–678.

— (2000). "HGLMs for analysis of correlated non-normal data." In *COMPSTAT: Proceedings in Computational Statistics 14th Symposium held in Utrecht, The Netherlands, 2000*, eds. J. G. Bethlehem and P. G. M. van der Heijden, 97–107. Utrecht, the Netherlands.

— (2001). "Modelling and analysing correlated non-normal data." *Statistical Modelling*, 1, 3–16.

Lehmann, E. (1999). *Elements of Large-Sample Theory*. New York, NY: Springer.

Lehmann, E. and Casella, G. (1998). *Theory of Point Estimation*. 2nd ed. New York, NY: Springer.

Neal, R. (2003). "Slice Sampling." *Annals of Statistics*, 31, 705–767.

Pourahmadi, M., Daniels, M. J., and Park, T. (2007). "Simultaneous modelling of the Cholesky decomposition of several covariance matrices." *Journal of Multivariate Analysis*, 98, 568–587.

Ravishanker, N. and Dey, D. K. (2002). *A First Course in Linear Model Theory*. Boca Raton, FL: Chapman and Hall/CRC.

Rue, H., Martino, S., and Chopin, N. (2009). "Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations." *Journal of the Royal Statistical Society, Series B*, 71, 319–392.

Sengupta, A., Cressie, N., Frey, R., and Kahn, B. (2012). "Statistical modeling of MODIS cloud data using the Spatial Random Effects model." In *Proceedings of the Joint Statistical Meetings*, 3111–3123. Alexandria, VA: American Statistical Association.

Shi, T. and Cressie, N. (2007). "Global statistical analysis of MISR aerosol data: A massive data product from NASA's Terra satellite." *Environmetrics*, 18, 665–680.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and Van Der Linde, A. (2002). "Bayesian measures of model complexity and fit." *Journal of the Royal Statistical Society, Series B*, 64, 583–616.

Torrieri, N. (2007). "America is changing, and so is the census: The American Community Survey." *American Statistician*, 61, 16–21.

Wikle, C. K. and Anderson, C. J. (2003). "limatological analysis of tornado report counts using a hierarchical Bayesian spatio-temporal model." *Journal of Geophysical Research-Atmospheres*, 108, 9005.

Wikle, C. K. and Cressie, N. (1999). "A dimension-reduced approach to space-time Kalman filtering." *Biometrika*, 86, 815–829.

Wikle, C. K. and Royle, J. A. (2005). "Dynamic design of ecological monitoring networks for nonGaussian spatio-temporal data." *Environmetrics*, 16, 507–522.

Wolpert, R. and Ickstadt, K. (1998). "Poisson/gamma random field models for spatial statistics." *Biometrika*, 85, 251–267.

Wu, G., Holan, S. H., and Wikle, C. K. (2013). "Hierarchical Bayesian Spatio-Temporal Conway-Maxwell Poisson Models with Dynamic Dispersion." *Journal of Agricultural, Biological, and Environmental Statistics*, 18, 335–356.

Yang, R. and Berger, J. (1994). "Estimation of a covariance matrix using the reference prior." *Annals of Statistics*, 22, 1195–1211.

Zhang, L., Guindani, M., and Vannucci, M. (2015). "Bayesian Models for fMRI Data Analysis." *Wiley Interdiscip Rev Comput Stat*, 7, 21–41.